AD-A174 225

Blackboard Systems

by

H. Penny Nii

Department of Computer Science

Stanford University
Stanford, CA 94305

DTIC
ELECTE
NOV 2 0 1986
S
E
D

86 11 19 125

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A174 225	

4. TITLE (and Subtitle)	5. TYPE OF REPORT & PERIOD COVERED
Blackboard Systems	technical
	6. PERFORMING ORG. REPORT NUMBER
	STAN-CS-86-1123

7. AUTHOR(s)	8. CONTRACT OR GRANT NUMBER(s)
H. Penny Nii	N0039-83-C-0136

9. PERFORMING ORGANIZATION NAME AND ADDRESS	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
Stanford Knowledge Systems Laboratory 701 Welch Road, Bldg C. Stanford University	

11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE
Defense Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209	June 1986
	13. NUMBER OF PAGES
	86

14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)	15. SECURITY CLASS. (of this report)
	Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release: distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

See next page.

DD FORM 1473
1 JAN 73

19. KEY WORDS (Continued)

20 ABSTRACT (Continued)

The first blackboard system was the HEARSAY-II speech understanding system [8], that evolved between 1971 and 1976. Subsequently, many systems have been built that have similar system organization and run-time behavior. The objectives of this document[1] are: (1) to define what is meant by "blackboard systems", and (2) to show the richness and diversity of blackboard system designs. The article begins with a discussion of the underlying concept behind all blackboard systems, the *blackboard model of problem solving*. In order to bridge the gap between a model and working systems, the *blackboard framework*, an extension of the basic blackboard model is introduced, including a detailed description of the model's components and their behavior. A model does not come into existence on its own and is usually an abstraction of many examples. In Section 2, the history of ideas is traced and the designs of some application systems that helped shape the blackboard model are detailed. We then describe and contrast existing blackboard systems. Blackboard systems can generally be divided into two categories; application and skeletal systems. In application systems the blackboard system components are integrated with the domain knowledge required to solve the problem at hand. Skeletal systems are devoid of domain knowledge, and, as the name implies, consist of the essential system components from which application systems can be built by the addition of knowledge and the specification of control (i.e. meta-knowledge). Application systems will be discussed in Section 3, and skeletal systems will be discussed elsewhere. In Section 3.6, we summarize the features of the application systems and in Section 4 present the author's perspective on the utility of the blackboard approach to problem solving and knowledge engineering.

Blackboard Systems

by

H. Penny Nii

Accession For	
NTIS GRA&I	X
DTIC TAB	☐
Unannounced	☐
Justification	

By		
Distribution/		
Availability Codes		
Dist	Avail and/or Special	
A-1		

KNOWLEDGE SYSTEMS LABORATORY
Departments of Medical and Computer Science
Stanford University
Stanford, California 94305

Table of Contents

List of Figures

Blackboard Systems

H. Penny Nii
Knowledge Systems Laboratory
Computer Science Department
Stanford University

The first blackboard system was the HEARSAY-II speech understanding system [8], that evolved between 1971 and 1976. Subsequently, many systems have been built that have similar system organization and run-time behavior. The objectives of this document[1] are: (1) to define what is meant by "blackboard systems", and (2) to show the richness and diversity of blackboard system designs. The article begins with a discussion of the underlying concept behind all blackboard systems, the *blackboard model of problem solving*. In order to bridge the gap between a model and working systems, the *blackboard framework*, an extension of the basic blackboard model is introduced, including a detailed description of the model's components and their behavior. A model does not come into existence on its own and is usually an abstraction of many examples. In Section 2, the history of ideas is traced and the designs of some application systems that helped shape the blackboard model are detailed. We then describe and contrast existing blackboard systems. Blackboard systems can generally be divided into two categories, application and skeletal systems. In application systems the blackboard system components are integrated with the domain knowledge required to solve the problem at hand. Skeletal systems are devoid of domain knowledge, and, as the name implies, consist of the essential system components from which application systems can be built by the addition of knowledge and the specification of control (i.e. meta-knowledge). Application systems will be discussed in Section 3, and skeletal systems will be discussed elsewhere. In Section 3.6, we summarize the features of the application systems and in Section 4 present the author's perspective on the utility of the blackboard approach to problem solving and knowledge engineering.

1. Blackboard Model of Problem Solving

Historically, the blackboard model arose from abstracting features of the HEARSAY-II speech-understanding system developed between 1971 and 1976. HEARSAY-II understood a spoken speech query about computer science abstracts stored in a database. It "understood" in the sense that it was able to respond to spoken commands and queries about the data-base. From an informal summary description of the HEARSAY-II program, the HASP system was designed

---

[1] This document is a part of a retrospective monograph on the AGE Project currently in preparation.

and implemented between 1973 and 1975. The domain[2] of HASP was ocean surveillance, and
its task[3] was the interpretation of continuous passive sonar data. HASP, as the second example
of a blackboard system, not only added credibility to the claim that a blackboard approach to
problem solving was general, but it also demonstrated that it could be abstracted into a robust
model of problem solving. Subsequently, many application programs have been implemented
whose solutions were formulated using the blackboard model. Because of the different
characteristics of the application problems and because the interpretation of the blackboard
model varied, the design of these programs differed considerably. However, the blackboard
model of problem solving has not undergone any substantial changes in the last ten years.

A *problem-solving model* is a scheme for organizing reasoning steps and domain knowledge to
construct a solution to a problem. For example, in a *backward-reasoning model*, problem
solving begins by reasoning backwards from a goal to be achieved towards an initial state
(data). More specifically, in a *rule-based backward-reasoning model* knowledge is organized
as "if-then" rules and modus ponens inference steps are applied to the rules from a goal rule
back to an "initial-state rule" (a rule that looks at the input data). An excellent example of
this approach to problem solving is the MYCIN program [45]. In a *forward-reasoning model*,
however, the inference steps are applied from an initial state toward a goal. The OPS system
exemplifies such a system [13]. In an *opportunistic-reasoning model*, pieces of knowledge are
applied either backward or forward at the most "opportune" time. Put another way, the central
issue of problem solving deals with the question of: "What pieces of knowledge should be
applied when and how?" A problem-solving model provides a conceptual framework for
organizing knowledge and a strategy for applying that knowledge.

The blackboard model of problem solving is a highly structured, special case of opportunistic
problem solving. In addition to opportunistic reasoning as a knowledge-application strategy,
the blackboard model prescribes the organization of the domain knowledge and all the input
and intermediate and partial solutions needed to solve the problem. We refer to all possible
partial and full solutions to a problem as its *solution space*.

In the blackboard model the solution space is organized into one or more application-

---

[2]*Domain* refers to a particular area of discourse, for example, chemistry.

[3]*Task* refers to a goal-oriented activity within the domain, for example, to analyze the molecular composition of a
compound.

dependent hierarchies.[4] Information at each level in the hierarchy represents partial solutions and is associated with a unique vocabulary that describes the information. The domain knowledge is partitioned into independent modules of knowledge that transform information on one level, possibly using information at other levels, of the hierarchy into information on the same or other levels. The knowledge modules perform the transformation using algorithmic procedures or heuristic rules that generate actual or hypothetical transformations. Opportunistic reasoning is applied within this overall organization of the solution space and task-specific knowledge; that is, which module of knowledge to apply is determined dynamically, one step at a time, resulting in the incremental generation of partial solutions. The choice of a knowledge module is based on the solution state (particularly, the latest additions and modifications to the data structure containing pieces of the solution) and on the existence of knowledge modules capable of improving the current state of the solution. At each step of knowledge application, either forward- or backward-reasoning methods may be applied.[5]

The blackboard model is a relatively complex problem-solving model prescribing the organization of knowledge and data and the problem-solving behavior within the overall organization. This section contains a description of the basic blackboard model. Variations and extensions will be discussed in subsequent sections.

1.1. The Blackboard Model
The blackboard model is usually described as consisting of three major components:

1. **The knowledge sources.** The knowledge needed to solve the problem is partitioned into *knowledge sources*, which are kept separate and independent.

2. **The blackboard data structure.** The problem-solving state data are kept in a global data base, the *blackboard*. Knowledge sources produce changes to the blackboard which lead incrementally to a solution to the problem. Communication and interaction among the knowledge sources take place solely through the blackboard.

3. **Control:** The knowledge sources *respond opportunistically to changes* in the

---

[4]The hierarchy may be an abstraction hierarchy, a part-of hierarchy, or any other type of hierarchy appropriate for solving the problem.

[5]There are various other ways of categorizing reasoning methods, for example, event driven, goal driven, model driven, expectation driven, and so forth. Without getting into the subtle differences between these methods, it is safe to say that any one of these methods can be applied at each step in the reasoning process.

blackboard.[6]

(--- Data flow)
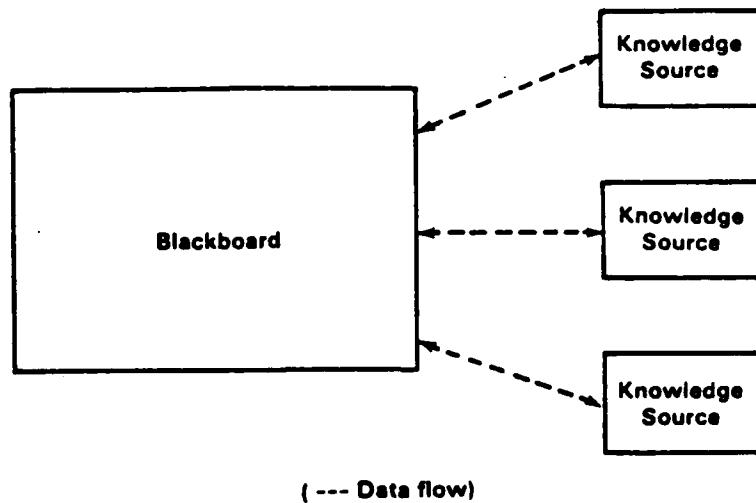
*There is a global database called the blackboard, and there are logically
independent sources of knowledge called the knowledge sources. The knowledge
sources respond to changes on the blackboard. Note that there is no control flow;
the knowledge sources are self-activating.*

Figure 1-1: The Blackboard Model

The difficulty with this description of the blackboard model is that it only outlines the
organizational principles. For those who want to build a blackboard system, the model does
not specify how it is to be realized as a computational entity, that is, the blackboard model is
a conceptual entity, not a computational specification. Given a problem to be solved, the
blackboard model provides enough guidelines for sketching a solution, but a sketch is a long
way from a working system. To design and build a system, a detailed model is needed. Before
moving on to adding details to the blackboard model, we explore the implied behavior of this
abstract model.

Let us consider a hypothetical problem of a group of people trying to put together a jigsaw
puzzle. Imagine a room with a large blackboard and around it a group of people each holding
over-size jigsaw pieces. We start with volunteers who put on the blackboard (assume it's
sticky) their most "promising" pieces. Each member of the group looks at his pieces and sees
if any of them fit into the pieces already on the blackboard. Those with the appropriate
pieces go up to the blackboard and update the evolving solution. The new updates cause other

---

[6]There is no control component specified in the blackboard model. The model merely specifies a general problem-
solving behavior. The actual locus of control can be in the knowledge sources, on the blackboard, in a seperate
module, or in some combination of the three. The need for a control component in blackboard systems is discussed
later.

pieces to fall into place, and other people go to the blackboard to add their pieces. It does not matter whether one person holds more pieces than another. The whole puzzle can be solved in complete silence; that is, there need be no direct communication among the group. Each person is self-activating, knowing when his pieces will contribute to the solution. No a priori established order exists for people to go up to the blackboard. The apparent cooperative behavior is mediated by the state of the solution on the blackboard. If one watches the task being performed, the solution is built incrementally (one piece at a time) and opportunistically (as an opportunity for adding a piece arises), as opposed to starting, say, systematically from the left top corner and trying each piece.
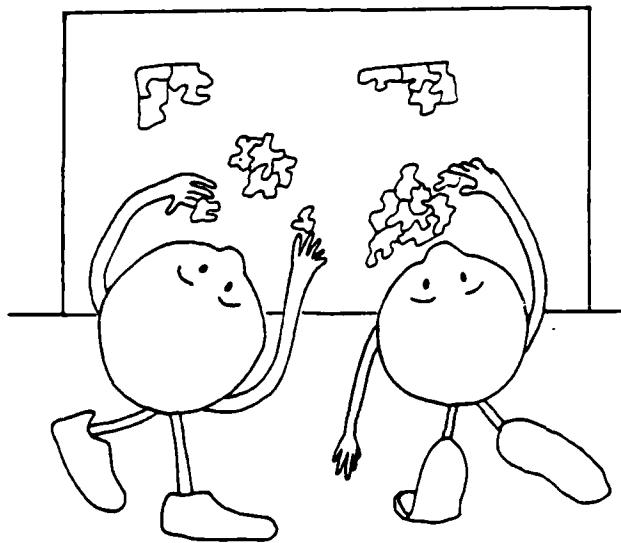
Figure 1-2: Solving Jigsaw Puzzles

This analogy illustrates quite well the blackboard problem-solving behavior implied in the model and is fine for a starter. Now, let's change the layout of the room in such a way that there is only one center aisle wide enough for one person to get through to the blackboard. Now, no more than one person can go up to the blackboard at one time, and a monitor is needed, someone who can see the group and can choose the order in which a person is to go up to the blackboard. The monitor can ask all people who have pieces to add to raise their hands. The monitor can then choose one person from those with their hands raised. To select one person, criteria for making the choice is needed, for example, a person who raises a hand first, a person with a piece that bridges two solution islands (that is, two clusters of completed pieces) and so forth. The monitor needs a strategy or a set of strategies for solving the puzzle. The monitor can choose a strategy before the puzzle solving begins or can develop strategies as the solution begins to unfold. In any case, it should be noted that the monitor has a broad

executive power. The monitor has so much power that the monitor could, for example, force the puzzle to be solved systematically from left to right; that is, the monitor has the power to violate one essential characteristic of the original blackboard model, that of opportunistic problem solving.

The last analogy, though slightly removed from the original model, is a useful one for computer programmers interested in building blackboard systems. Given the serial nature of most current computers, the conceptual distance between the model and a running blackboard system is a bit far, and the mapping from the model to a system is prone to misinterpretation. By adding the constraint that solution building physically occur one step at a time in some order determined by the monitor (when multiple steps are possible and desirable), the blackboard model is brought closer to the realities inherent in serial computing environments.[7]

Although the elaborate analogy to jigsaw puzzle solving gives us additional clues to the nature of the behavior of blackboard systems, it is not a very good example for illustrating the organization of the blackboard or for the partitioning of appropriate knowledge into knowledge sources. To illustrate these aspects of the model, we need another example. This time let us consider another hypothetical problem, that of finding koalas in a eucalyptus forest (see Figure 1-3).

Imagine yourself in Australia. One of the musts if you are a tourist is to go and look for koalas in their natural habitat. So, you go to a koala preserve and start looking for them among the branches of the eucalyptus trees. You find none. You know that they are rather small, grayish creatures which look like bears.[8] The forest is dense, however, and the combination of rustling leaves and the sunlight reflecting on the leaves adds to the difficulty of finding these creatures, whose coloring is similar to their environment.[9] You finally give up and ask a ranger how you can find them. He gives you the following story about koalas: "Koalas usually live in groups and seasonally migrate to different parts of the forest, but they

---

[7]The serialization of the blackboard model is useful only because we tend to work on uniprocessor computers. We are currently conducting research on concurrent problem-solving methods. A starting point for the work is the pure blackboard model. One can see, at least conceptually, much parallelism inherent in the model. The problem is how to convert the model into an operational system that can take advantage of many (100s to 1000s) processor-memory pairs.

[8]More details at this descriptive level would be considered factual knowledge and can be used as a part of a prototypical model of koalas.

[9]The signal-to-noise ratio is low.

should be around the northwest area of the preserve now. They usually sit on the crook of branches and move up and down the tree during the day to get just the right amount of sun.[10] If you are not sure whether you have spotted one, watch it for a while; it will move around, though slowly.[11]" Armed with the new knowledge, you go back to the forest with a visual image of exactly where and what to look for. You focus your eyes at about 30 feet with no luck, but you try again, and this time focus your eyes at 50 feet, and suddenly you do find one. Not only one, but a whole colony of them.[12]

Figure 1-3:   Finding Koalas

Let's consider one way of formulating this problem along the lines of the blackboard model. Many kinds of knowledge can be brought to bear on the problem: the color and shape of koalas, the general color and texture of the environment (the noise characteristics), the behavior of the koalas, effects of season and time of the day, and so on. Some of the

---

[10]This is knowledge about the prototypical behavior pattern of koalas. The ranger suggests a highly model-driven approach to finding them.

[11]This is a method of detection as well as confirmation.

[12]This koala problem has a long history. It was invented by Ed Feigenbaum (after his trip to Australia) and myself in 1974, during the time when we were not allowed to write about the HASP project. The primary objective of this example was to illustrate the power of model-directed reasoning in interpreting noisy data. A paper was written about it but has been collecting dust, a victim of our distaste for writing about hypothetical problems. I resurrect it here because it's hard to come up with a good example that does not require specialized domain knowledge.

knowledge can be found in books, such as a <u>Handbook of Koala Sizes and Color</u> or <u>Geography of the Forest</u>. Some knowledge is informal -- the most likely places to find koalas at any given time or their favorite resting places. How can these diverse sources of knowledge be used effectively? First, we need to decide what constitutes a solution to the problem. Then, we can consider what kinds of information are in the data, what can be inferred from them, and what knowledge might be brought to bear to achieve the goal of finding the koalas.

Think of the solution to this problem as a set of markings on a series of snapshots of the forest. The markings might say, "This is certainly a koala because it has a head, body, and limbs and because it has changed its position since the last snapshot;" or "This might be a koala, because it has a blob that looks like a head;" or "These might be koalas because they are close to the one we know is a koala and the blobs could be heads, legs, or torsos." The important characteristics of the solution are that the solution consists of bits and pieces of information, and it is a reasoned solution with supporting evidence and supporting lines of reasoning.

Having decided that the solution would consist of partial and hypothetical identifications, as well as complete identifications constructed from partial ones, we need a solution-space organization that can hold descriptions of bits and pieces of the koalas. One such descriptive framework is a part-of hierarchy. For each koala, the highest level of description is the koala itself, which is described on the next level by head and body; the head is described on the next level by ears, nose, and eyes; the body is described by torso, legs, and arms; and so on. At each level, there are descriptors appropriate for that level; size, gender, and height on the koala level, for example. Each primitive body part is described on the lower levels in terms of geometric features, such as shapes and line segments. Each shape has color and texture associated with it as well as its geometric descriptions (see Figure 1-4). In order to identify a part of the snapshot as a koala, we need to mark the picture with line segments and regions. The regions and pieces of lines must eventually be combined, or synthesized, in such a way that the description of the constructed object can be construed as some of the parts of a koala or a koala itself. For example, a small, black circular blob could be an eye, but it must be surrounded by a bigger, lighter blob that might be a head. The more pieces of information one can find that fit the koala description, the more confident we can be. In addition to the body parts that support the existence of a koala, if the hypothesized koala is at about 30 to 50 feet above ground, we would be more confident than if we found the same object at 5 feet.

The knowledge needed to fill in the koala descriptions falls into place with the decision to organize the solution space as a part-of abstraction hierarchy. We would need a color specialist, a shape specialist, a body-part specialist, a habitat specialist, and so forth. No one source of knowledge can solve the problem; the solution to the problem depends on the
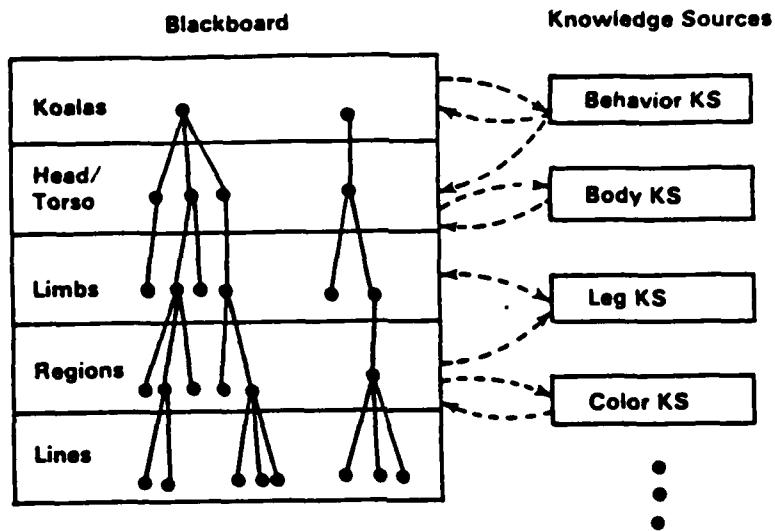
The koalas in the scene are described as a part-of hierarchy. Specialist knowledge modules contribute information about what they "see" to help in the search for koalas.

Figure 1-4: Koalas: Blackboard Structure and Knowledge Sources

combined contributions of many specialists. The knowledge held by these specialists is logically independent. Thus, a color specialist can determine the color of a region without knowing how the shape specialist determined the shape of the region. However, the solution of the problem is dependent on both of them. The torso specialist does not have to know whether the arm specialist checked if an arm had paws or not (the torso specialist probably doesn't even know about paws), but each specialist must rely on the other specialists to supply the information each needs. Cooperation is achieved by assuming that whatever information is needed is supplied by someone else.

The jigsaw puzzle and the koala problems illustrate the organization of information on the blackboard database, the partitioning of domain knowledge into specialized sources of knowledge, and some of the characteristic problem-solving behavior associated with the blackboard model.[13] Neither of these, however, answers the questions of how the knowledge is to be represented, or of what the mechanisms are for determining and activating appropriate knowledge. As mentioned earlier, problem-solving models are conceptual frameworks for formulating solutions to problems. The models do not address the details of designing and

---

[13]As in the jigsaw problem, the problem-solving behavior in the koala problem would be opportunistic. As new pieces of evidence are found and new hypotheses generated, appropriate knowledge sources analyze them and create new hypotheses.

building operational systems. How a piece of knowledge is represented, as rules, objects, or procedures, is an engineering decision. It involves such pragmatic considerations as "naturalness," availability of a knowledge representation language, and the skill of the implementers, to name but a few.[14] What control mechanisms are needed depends on the complexity and the nature of the application task. We can, however, attempt to narrow the gap between the model and operational systems. Now, the blackboard model is extended by adding more details to the three primary components in terms of their structures, functions, and behaviors.

1.2. The Blackboard Framework

Applications are implemented with different combinations of knowledge representations, reasoning schemes, and control mechanisms. The variability in the design of blackboard systems is due to many factors, the most influential one being the nature of the application problem itself. It can be seen, however, that blackboard architectures which underly application programs have many similar features and constructs. (Some of the better known applications are discussed in Section 3.) The *blackboard framework* is created by abstracting these constructs.[15] The blackboard framework, therefore, contains descriptions of the blackboard system components that are grounded in actual computational constructs. The purpose of the framework is to provide design guidelines appropriate for blackboard systems in a serial-computing environment.[16] Figure 1-5 shows some modifications to Figure 1-1 to reflect the addition of system-oriented details.

1. **The knowledge sources:** The domain knowledge needed to solve a problem is partitioned into *knowledge sources* that are kept separate and independent.

> *The objective of each knowledge source is to contribute information that will lead to a solution to the problem.* A knowledge source takes a set of current information

---

[14]The blackboard model does not preclude the use of human knowledge sources. Interesting interactive and symbiotic expert systems can be built by integrating human expertise during run time.

[15]There is an implicit assumption that systems can be described at various levels of abstraction. Thus, the description of the framework is more detailed than the model and less detailed than a specification (a description from which a system can be built). Here, they are called the model, framework, and specification levels.

[16]One can view the blackboard framework as a prescriptive model; that is, it prescribes what must be in a blackboard system. However, it must be kept in mind that application problems often demand extensions to the framework, as can be seen in the examples in Section 3.

on the blackboard and updates it as encoded in its specialized knowledge.

The knowledge sources are represented as procedures, sets of rules, or logic assertions. To date most of the knowledge sources have been represented as either procedures or as sets of rules. However, systems that deal with signal processing either make liberal use of procedures in their rules, or use both rule sets and procedurally encoded knowledge sources.

The knowledge sources modify only the blackboard or control data structures (that also might be on the blackboard), and only the knowledge sources modify the blackboard. All modifications to the solution state are explicit and visible.

Each knowledge source is responsible for knowing the conditions under which it can contribute to a solution. Each knowledge source has *preconditions* that indicate the condition on the blackboard which must exist before the body of the knowledge source is *activated.*[17]

2. **The blackboard data structure:** The problem-solving state data are kept in a global database, the *blackboard.* Knowledge sources produce changes to the blackboard that lead incrementally to a solution, or a set of acceptable solutions, to the problem. Interaction among the knowledge sources takes place solely through changes on the blackboard.

The purpose of the blackboard is to hold computational and solution-state data needed by and produced by the knowledge sources. The knowledge sources use the blackboard data to interact with each other indirectly.

The blackboard consists of objects from the solution space. These objects can be input data, partial solutions, alternatives, and final solutions (and, possibly, control data).

The objects on the blackboard are hierarchically organized into levels of analysis. Information associated with objects (that is, their properties) on one level serves as input to a set of knowledge sources, which, in turn, place new information on the same or other levels.

The objects and their properties define the vocabulary of the solution space. The properties are represented as attribute-value pairs. Each level uses a distinct subset

---

[17]One can view a knowledge source as a large rule. The major difference between a rule and a knowledge source is the grain size of the knowledge each holds. The condition part of this large rule is called the *knowledge source precondition,* and the action part is called the *knowledge source body.*

(— Control flow; --- Data flow)
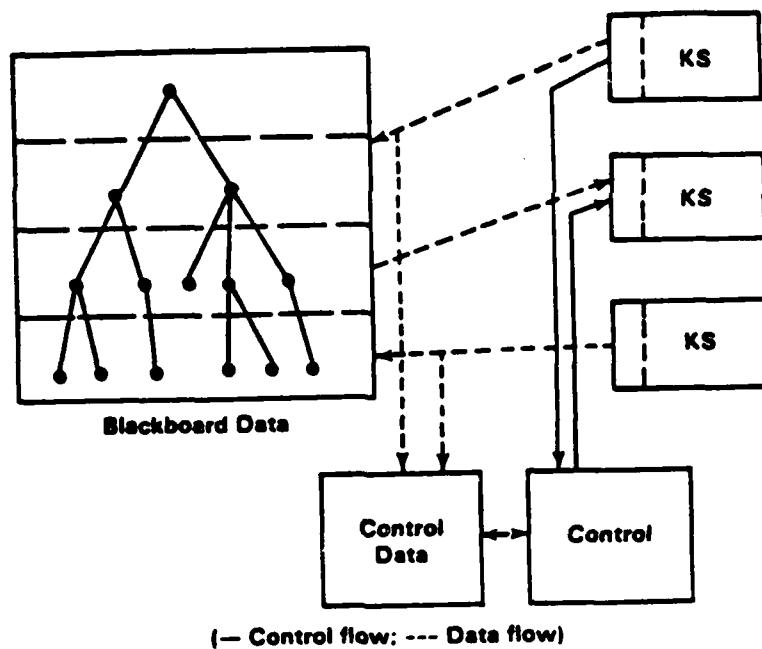
The data on the blackboard are hierarchically organized. The knowledge sources are logically independent, self-selecting modules. Only the knowledge sources are allowed to make changes to the blackboard. Based on the latest changes to the information on the blackboard, a control module selects and executes the next knowledge source.

Figure 1-5: The Blackboard Framework

of the vocabulary.[18]

The relationships between the objects are denoted by named links.[19] The relationship can be between objects on different levels, such as "part-of" or "in-support-of," or between objects on the same level, such as "next-to" or "follows."

The blackboard can have multiple blackboard panels; that is, a solution space can be partitioned into multiple hierarchies.[20]

3. **Control:** The knowledge sources *respond opportunistically to changes* on the blackboard.

---

[18]Many times, the names of the attributes on different levels are the same, for example "type." Often these are shorthand notations for "type-of-x-object" or "type-of-y-object." Some times they are duplications of the same attribute for convenience sake.

[19]A relationship is a special kind of property.

[20]This feature was first used in the CRYSALIS system. The rationale for introducing multiple panels is discussed in Section 3.3.

There is a set of control modules that monitor the changes on the blackboard and decide what actions to take next.

Various kinds of information are made globally available to the control modules. The information can be on the blackboard or kept separately. The control information is used by the control modules to determine the *focus of attention.*

The focus of attention indicates the next thing to be processed. The focus of attention can be either the knowledge sources (that is, which knowledge sources to activate next) or the blackboard objects (i.e., which solution islands to pursue next), or a combination of both (i.e., which knowledge sources to apply to which objects).[21]

The solution is built one step at a time. Any type of reasoning step (data driven, goal driven, model driven, and so on) can be applied at each stage of solution formation. As a result, the sequence of knowledge source invocation is dynamic and opportunistic rather than fixed and preprogrammed.

Pieces of problem-solving activities occur in the following iterative sequence:

1. A knowledge source makes change(s) to blackboard object(s). As these changes are made, a record is kept in a global data structure that holds the control information.

2. Each knowledge source indicates the contribution it can make to the new solution state. (This can be defined *a priori* for an application, or dynamically determined.)

3. Using the information from points 1 and 2 a control module selects a focus of attention.

4. Depending on the information contained in the focus of attention, an appropriate control module prepares it for execution as follows:

   a. If the focus of attention is a knowledge source, then a blackboard object (or sometimes, a set of blackboard objects) is chosen to serve as the context of its invocation (knowledge-scheduling approach).

   b. If the focus of attention is a blackboard object, then a knowledge source is chosen which will process that object (event-scheduling approach).

   c. If the focus of attention is a knowledge source and an object, then that knowledged source is ready for execution. The knowledge source is executed together with the context, thus described.

Criteria are provided to determine when to terminate the process. Usually, one of the

---

[21]Any given system usually employs one of the three approaches, not all.

knowledge sources indicates when the problem-solving process is terminated, either because an acceptable solution has been found or because the system cannot continue further for lack of knowledge or data.

Problem-Solving Behavior and Knowledge Application

The problem-solving behavior of a system is determined by the knowledge-application strategy encoded in the control modules. The choice of the most appropriate knowledge-application strategy is dependent on the characteristics of the application task and on the quality and quantity of domain knowledge relevant to the task.[22] Basically, the acts of choosing a particular blackboard region and choosing a particular knowledge source to operate on that region determine the problem-solving behavior. Generally, a knowledge source uses information on one level as its input and produces output information on another level. Thus, if the input level of a particular knowledge source is on the level lower (closer to data) than its output level, then the application of this knowledge source is an application of bottom-up, forward reasoning.

Conversely, a commitment to a particular type of reasoning step is a commitment to a particular knowledge-application method. For example, if we are interested in applying a data-directed, forward-reasoning step, then we would select a knowledge source whose input level is lower than its output level. If we are interested in goal-directed reasoning, we would select a knowledge source that put information needed to satisfy a goal on a lower level. Using the constructs in the control component one can make any type of reasoning step happen at each step of knowledge application.[23]

How a piece of knowledge is stated often presupposes how it is to be used. Given a piece of knowledge about a relationship between information on two levels, that knowledge can be

---

[22]It might be said that this is a hedge, that there should be a knowledge-application strategy or a set of strategies built into the framework to reflect different problem-solving behaviors. It is precisely this lack of doctrine that makes the blackboard framework powerful and useful. If an application task calls for two forward-reasoning steps followed by three backward-reasoning steps at some particular point, the framework allows for this. This is not to say that a system with built-in strategies cannot be designed and built. If there is a knowledge-application strategy "generic" to a class of applications, then it might be worthwhile to build a skeletal system with that particular strategy.

[23]The control component of the framework is extensible in many directions. In the BB-1 system [18], the control problem is viewed as a planning problem. Knowledge sources are applied according to a problem-solving plan in affect. The creation of a problem-solving plan is treated as another problem to be solved using the blackboard approach.

expressed in top-down or bottom-up application forms. These can further be refined. The top-down form can be written as a goal, an expectation, or as an abstract model of the lower-level information. For example, a piece of knowledge can be expressed as a conjunction of information on a lower level needed to generate a hypothesis at a higher level (a goal), Or, it can be expressed as information on a lower level needed to confirm a hypothesis at a higher level (an expectation), and so on. The framework does not presuppose nor does it prescribe the knowledge-application, or reasoning, methods. It merely provides constructs within which any reasoning methods can be used. Many interesting problem-solving behaviors have been implemented using these constructs, some of them are discussed in Section 3.

1.3. Perspectives

The organizational underpinnings of blackboard systems have been the primary focus. The blackboard framework is a system-oriented interpretation of the blackboard model. It is a mechanistic formulation intended to serve as a foundation for system specifications. In problem-solving programs, we are usually interested in their performance and problem-solving behavior, not their organization. We have found, however, that some classes of complex problems become manageable when they are formulated along the lines of the blackborad model. Also, interesting problem-solving behavior can be programmed using the blackboard framework as a foundation. Even though the blackboard framework still falls short of being a computational specification, given an application task and the necessary knowledge, it provides enough information so that a suitable blackboard system can be designed, specified, and built. Some examples of complex problems with interesting problem solving behavior are discussed in Section 3. The examples show that new constructs can be added to the blackboard framework as the application problems demand, without violating the guidelines contained in it.[24]

There are other perspectives on the blackboard model. The blackboard model is sometimes viewed as a model of general problem solving [16]. It has been used to structure cognitive

---

[24]What about statements such as: "Fortran Common is a blackboard." or "Object-uriented systems are blackboard systems." All I can say is *The potential for a thing is not that thing itself*. With some effort, one can design and build a blackboard system in Fortran, and the Common area is a good candidate for storing blackboard data. However, one also needs to design knowledge sources that are self-selecting and self-contained and control modules that determine the focus of attention and manage knowledge source application. The blackboard framework is a problem solving framework. It is not a programming language, although an instance of the Framework can have a blackboard language associated with it. It is not a knowledge representation language, although one can use any knowledge representation language for the knowledge sources and the blackboard. Why can't I get away with placing a hunk of ground beef, a can of tomato sauce, a box of spaghetti, and bottles of seasoning in a pile and call the pile a spaghetti dinner or, better yet, linguine a la pesta rosa? It would certainly simplify my life.

models [28], [41], [20]; the OPM system (described in Section 3.5.1) simulates the human planning process. Sometimes the blackboard model is used as an organizing principle for large, complex systems built by many programmers. The ALVan project [49] takes this approach.

1.4. Summary

The basic approach to problem solving in the blackboard framework is to divide the problem into loosely coupled subtasks. These subtasks roughly correspond to areas of specialization within the task (for example, there are human specialists for the subtasks). For a particular application, the designer defines a solution space and knowledge needed to find the solution. The solution space is divided into analysis levels of partial or intermediate solutions, and the knowledge is divided into specialized knowledge sources that perform the subtasks. The information on the analysis levels is globally accessible on the blackboard, making it a medium of interaction between the knowledge sources. Generally, a knowledge source uses information on one level of analysis as its input and produces output information on another level. The decision to employ a particular knowledge source is made dynamically using the latest information contained in the blackboard data structure (the current solution state). This particular approach to problem decomposition and knowledge application is very flexible and works well in diverse application domains. One caveat, however: How the problem is partitioned into subproblems makes a great deal of difference to the clarity of the approach, the speed with which solutions are found, the resources required, and even the ability to solve the problem at all.

In order to discuss the details of various blackboard systems, it is helpful to trace the intellectual history of the blackboard concepts. Aside from being interesting in itself, it explains the origins of ideas and reasons for some of the differences between blackboard system designs. The reasons often have no rational basis but have roots in the "cultural" differences between the research laboratories that were involved in the early history of blackboard systems.

2. Evolution of Blackboard Architectures

"Metaphorically we can think of a set of workers, all looking at the same blackboard:
each is able to read everything that is on it, and to judge when he has something
worthwhile to add to it. This conception is just that of Selfridge's Pandemonium
[42]: a set of demons, each independently looking at the total situation and
shrieking in proportion to what they see that fits their natures..."
[*Allen Newell 1962*]

2.1. Prehistory

The above quotation is the first reference to the term *blackboard* in the AI literature. Newell
was concerned with the organizational problems of programs that existed at the time (for
example, checker-playing programs, chess-playing programs, theorem-proving programs), most
of which were organized along a generate-and-test search model[25] [31]. (See Feigenbaum and
Feldman 1963 [10] for a collection of articles describing some of these programs.) The major
difficulty in these programs was rigidity. He notes:

"...a program can operate only in terms of what it knows. This knowledge can come
from only two sources. It can come from assumptions [or] it can come from
executing processes ... either by direct modification of the data structure or by testing
... but executing processes take time and space [whereas] assumed information does
not have to be stored or generated. Therefore the temptation in creating efficient
programs is always to minimize the amount of generated information, and hence to
maximize the amount of stipulated information. It is the latter that underlies most
of the rigidities."

In one example, Newell discusses an organization to synthesize complex processes by means of
sequential flow of control and hierarchically organized, closed subroutines. Even though this
organization had many advantages (isolation of tasks, space saving by coding nearly identical
tasks once, and so on), it also had difficulties. First, conventions required for communication
among the subroutines often forced the subroutines to work with impoverished information.
Second, the ordered subroutine calls fostered the need for doing things sequentially. Third, and
most importantly, it encouraged programmers to think of the total program in terms of only
one thing going on at a time. However, in problem solving there are often many possible
things to be processed at any given time (for example, exploring various branches of a search
tree), and relatively weak and scattered information is necessary to guide the exploration for a
solution (for example, observations noticed while going down one branch of a search tree could
be used when going down another branch). The primary difficulties with this organiztion,

---

[25]"Generate" is a process that produces each element of the solution space one at a time. The "test" process
determines if the generated element satisfied some conditions of predicates in the task domain.

then, were inflexible control and restricted data accessibility. It is within this context that Newell notes that the difficulties "might be alleviated by maintaining the isolation of routines, but allowing all the subroutines to make use of a common data structure." He uses the blackboard metaphor to describe such a system.

The blackboard solution proposed by Newell eventually became the production system [32], which in turn led to the development of the OPS system [13]. In OPS, the "subroutines" are represented as *condition-action rules*,[26] and the data are globally available in the *working memory*. One of the many "shrieking demons" (those rules whose "condition sides" are satisfied) is selected through a *conflict-resolution* process. The conflict resolution process emulates the selection of one of the loudest demons, for example, one that addresses the most specific situation. OPS does reflect the blackboard concept as stated by Newell and provides for flexibility of control and global accessibility to data. However, the blackboard systems as we know them today took a slightly more circuitous route before coming into being.

In a paper first published in 1966 (later published in [47]), Simon mentions the term blackboard in a slightly different context from Newell. The discussion is within the framework of an information processing theory about discovery and incubation of ideas:

> "In the typical organization of problem-solving program, the solution efforts are guided and controlled by a hierarchy or tree of goals and subgoals. Thus, the subject starts out with the goal of solving the original problem. In trying to reach this goal, he generates a subgoal. If the subgoal is achieved, he may then turn to the now-modified original goal. If difficulties arise in achieving the subgoal, sub-subgoals may be created to deal with them ... we would specify that the goal tree be held in some kind of temporary memory, since it is a dynamic structure, whose function is to guide search, and it is not needed when the problem solution has been found ... In addition, the problem solver is noticing various features of the problem environment and is storing some of these in memory ... What use is made of [a feature] at the time it is noted depends on what subgoal is directing attention at that moment ... over the longer run, this information influences the growth of the subgoal tree ... I will call the information about the task environment that is noticed in the course of problem solution and fixated in permanent (or relatively long-term) memory the 'blackboard' ..."

Although Newell's and Simon's concerns appear within different contexts, the problem-solving method they were using was the goal-directed, generate-and-test search method. They encountered two common difficulties: the need for previously generated information during

---

[26]See [6] for an overview of production systems.

problem solving and flexible control. It was Simon who proposed the blackboard ideas to Raj
Reddy and Lee Erman for the HEARSAY project.[27]

Although the blackboard metaphor was suggested by Simon to the HEARSAY designers, the
final design of the system, as might be expected, evolved out of the needs of the speech-
understanding task. Such system characteristics as hierarchically organized analysis levels on
the blackboard and opportunistic reasoning, which we now accept as integral parts of
blackboard systems, were derived from needs and constraints that were different from Newell's
and Simon's. One of the key notions attributable to the speech-understanding problem was the
notion of the blackboard partitioned into analysis levels. This is a method of using and
integrating different "vocabularies," as mentioned earlier, in problem solving. In most
problem-solving programs of the time, such as game-playing and theorem-proving programs,
the problem space had a homogeneous vocabulary. In the speech-understanding problem, there
was a need to integrate concepts and vocabularies used in describing grammars, words, phones,
and so on.

There are two interesting observations to be made from early history. First, the early allusions
to a blackboard are closely tied to search methodologies, and, not surprisingly, the use of
generate-and-test search is evident in HEARSAY-II. Second, although the HEARSAY-II
blackboard system was designed independently from the OPS system, there are, as we might
expect, some conceptual similarities. For example, the *scheduler* in HEARSAY-II is
philosophically and functionally very similar to the *conflict-resolution* module in OPS, which,
in turn. is a way of selecting one of the shrieking demons.

The HASP system, which has its own intellectual history, does not focus so much on search
techniques as on knowledge-application techniques. Put another way, HASP was built in a
culture that had a tradition of using problem-solving approaches that focused on applying large
amounts of situation-specific knowledge rather than in applying a weak method (generate-and-
test) using general knowledge about the task.[28] The methodology used to select and apply
knowledge in HASP is, therefore, quite different philosophically from the one reflected in the

---

[27]These historical notes are communications from Herbert Simon.

[28]Most OPS expert systems use strong knowledge, but this came about later (ca 1979)

HEARSAY-II scheduler.[29] These and other differences are elaborated on in Sections 2.2 and 2.3. Next, another branch of a history that influenced the design of HEARSAY-II is examined, the speech-understanding task.

2.2. The HEARSAY Project

Although a blackboard concept was documented in AI literature as early as 1962 by Newell, it was implemented as a system a decade later by people working on a speech-understanding project. The first article on the HEARSAY system appeared in the *IEEE Transactions on Audio and Electroacoustics* in 1973 [38].[30] There, the authors described the limitations of extant speech-recognition systems and proposed a model that would overcome the limitations. To summarize, the article stated that although the importance of context, syntax, semantics, and phonological rules in the recognition of speech was accepted, no system had been built that incorporated these ill-defined sources of knowledge. At the same time, the authors' previous work indicated (1) that the limitation of syntax-directed methods of parsing from left to right had to be overcome; (2) that parsing should proceed both forward and backward from anchor points; and (3) that because of the lack of feedback in a simple part-of hierarchical structure, the magnitude of errors on the lower level propagated multiplicatively up the hierarchy; that is, minor errors in the signal level, for example, became major errors on a sentence level.

The system architecture described in the Reddy article, later to be known as the HEARSAY-I architecture, was based on a model that addressed the following requirements: (1) the contribution of each source of knowledge (syntax, semantics, context, and so on) to the recognition of speech had to be measurable; (2) the absence of one or more knowledge sources should not have a crippling effect on the overall performance; (3) more knowledge sources should improve the performance; (4) the system must permit graceful error recovery; (5) changes in performance requirements, such as increased vocabulary size or modifications to the syntax or semantics, should not require major modifications to the model. The functional diagram of the HEARSAY-I architecture is shown in Figure 2-1, and its behavior is summarized as follows:

---

[29]There is no denying that there are cultural differences in the AI laboratories; they foster different styles, methods, and lines of research. Whatever the research topic, the intellectual effort tends to follow the line of least resistance and adopt the styles and methods at the researcher's own laboratory. Thus, the work of Newell and Simon on general problem solving has a great deal of influence on much of the work at Carnegie-Mellon University, whereas, the work of Feigenbaum and Buchanan on applications of domain-specific knowledge influences the work at their Stanford University laboratory.

[30]The manuscript was delivered to IEEE on April 30, 1972.

"The EAR module accepts speech input, extracts parameters, and performs some
preliminary segmentation, feature extraction, and labeling, generating a "partial
symbolic utterance description." The recognition overlord (ROVER) controls the
recognition process and coordinates the hypothesis generation and verification phases
of various cooperating parallel processes. The TASK provides the interface between
the task being performed and the speech recognition and generation (SPEAK-EASY)
parts of the system. The system overlord (SOL) provides the overall control for the
system."
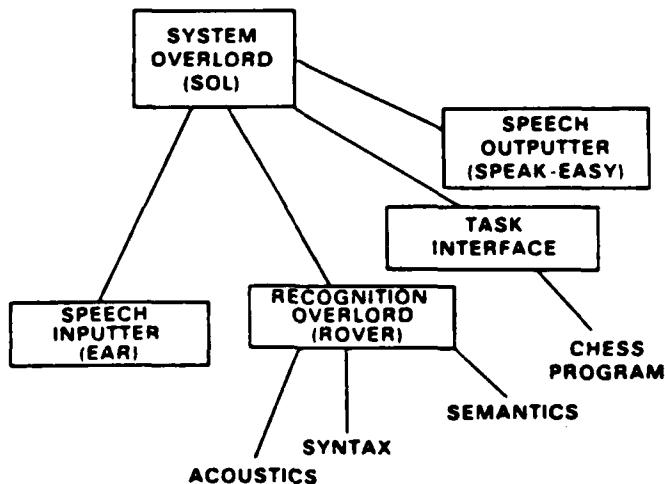
Figure 2-1:   Overview of the HEARSAY-I System -- from [39]

From Figure 2-2 which illustrates the recognition process, one can glean the beginnings of an
organization of a blackboard system. Note how the overlord (ROVER) controlled the
invocation of activities. The beginnings of the scheduler, as well as the knowledge sources are
apparant, as they became incorporated in HEARSAY-II.

"Since the different recognizers are independent, the recognition overlord needs to
synchronize the hypothesis generation and verification phases of various processes.
... Several strategies are available for deciding which subset of the processes
generates the hypotheses and which verify. At present this is done by polling the
processes to decide which process is most confident about generating the correct
hypothesis. In voice chess, [The task domain for HEARSAY-I was chess moves.],
where the semantic source of knowledge is dominant, that module usually generates
the hypotheses. These are then verified by the syntactic and acoustic recognizers.
However, when robust acoustic cues are present in the incoming utterance, the roles
are reversed with the acoustic recognizer generating the hypotheses."
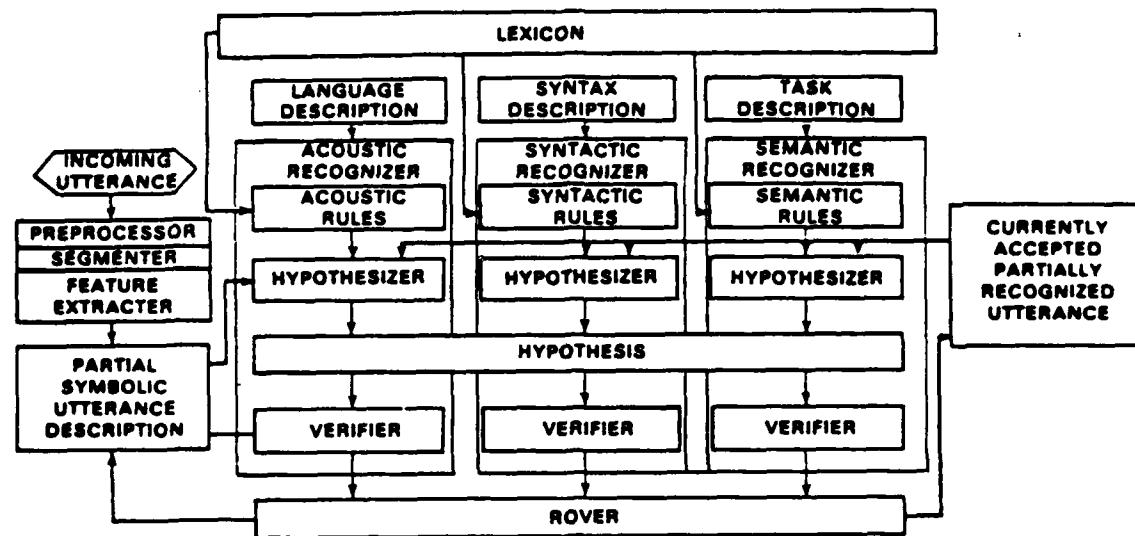
Figure 2-2:   Details of the Recognition Process -- from [38]

"Knowledge sources are activated in a lock-step sequence consisting of three phases: poll, hypothesize, and test," [24]. During the polling phase, the overlord queries the knowledge sources to determine which ones have something to contribute to that region of the sentence hypothesis which is "in focus" and with what level of "confidence."[31] In the hypothesizing phase, the most promising knowledge source is activated to make its contribution. Finally, in the testing phase, knowledge sources evaluate the new hypotheses.

Some of the difficulties encountered in HEARSAY-I can be attributed to the way in which the solution to the application task was formulated, and other difficulties arose from the design of the system. The problem was formulated to use the hypothesize-and-test paradigm only on the word level, that is, the blackboard only contained a description at the word level. This meant that all communication among the knowledge sources was limited to sharing information at the word level. This formulation caused two major difficulties. First, it becomes difficult to add non-word knowledge sources and to evaluate their contributions. Second, the inability to share information contributed by non-word knowledge sources caused the information to be recomputed by each knowledge source that needed it. In other words, the difficulty lay in trying to force the use of a single vocabulary (that is, from the word level) when multiple vocabularies (for example, on the acoustic level) were needed.

---

[31]The poll portion of the poll, hypothesize, and test is also very characteristic of OPS and HEARSAY-II. This construct takes on a totally different form in HASP and other subsequent systems. ·

The architectural weaknesses of HEARSAY-I, as stated by its designers, lay in (1) the lock-step control sequence that limited "parallelism,"[32] (2) the lack of provision to express relationships among alternative sentence hypotheses, and (3) the built-in problem-solving strategy that made modifications awkward and comparisons of different strategies impossible. [24] To overcome these difficulties, information (in the multiple vocabularies needed to understand utterances) used by all the knowledge sources was uniformly represented and made globally accessible on the blackboard in HEARSAY-II. In addition, a scheduler dynamically selected and activated the appropriate knowledge sources. (In Section 3.1 the design of the HEARSAY-II system is described in detail.)

During the time that HEARSAY-II was being developed, the staff of the HASP project was looking for an approach to solve its application problem. The search for a new methodology came about because the plan-generate-and-test problem-solving method that was successful for interpreting mass-spectrometry data in the DENDRAL program [27] was found to be inappropriate for the problem of interpreting passive sonar signals. In the history of blackboard systems, HASP represents a branching point in the philosophy underlying the design of blackboard systems. Generally, later systems can be thought of as modifications of, or extensions to either the HEARSAY-like or HASP-like designs.

2.3. The HASP Project
The task of HASP was to interpret continuous sonar signals passively collected by hydrophone arrays monitoring an area of the ocean. Signals are received from multiple arrays, with each array consisting of multiple hydrophones. Each array has some directional resolution. Imagine a large room full of plotters, each recording digitized signals from the hydrophones. Now, imagine an analyst going from one plotter to the next trying to discern what each one is hearing, and then integrating the information from all the plots in order to discern the current activity in the region under surveillance. This interpretation and analysis activity goes on continuously day in and day out. The primary objective of this activity is to detect enemy submarines. The objective of the HASP project was to write a program that "emulated" the human analysts, that is, to incorporate, in a computer program, the expertise of the analysts, especially their ability to detect submarines.[33] The HASP problem was chosen to work on because it appeared to be similar to the DENDRAL problem, a signal interpretation problem

---

[32]The term parallelism was used quite early in the project even though at that time the system ran on uniprocessors. Later (ca. 1976), experiments with parallel executions were conducted on the C.mmp system. [12].

[33]This was in 1973 before the term expert system was coined. The only expert system in existence at the time was DENDRAL, and MYCIN was on its way.

for which there were experts who could do the job. The system designers were confident that the problem-solving approach taken in DENDRAL would work for HASP. What was DENDRAL's task, and what was its approach? To quote from [11], the task was

> to enumerate plausible structures (atom-bond graphs) for organic molecules, given two kinds of information: analytic instrument data from a mass spectrometer and a nuclear magnetic resonance spectrometer; and user-supplied constraints on the answers, derived from any other source of knowledge (instrumental or contextual) available to the user.
>
> DENDRAL's inference procedure is a heuristic search that takes place in three stages, without feedback: *plan-generate-and-test*.
>
> *Generate* is a generation process for plausible structures. Its foundation is a combinatorial algorithm that can produce all the topologically legal candidate structures. Constraints supplied by the user or by the *Plan* process prune and steer the generation to produce the plausible set and not the enormous legal set.
>
> *Test* refines the evaluation of plausibility, discarding less worthy candidates and rank-ordering the remainder for examination by the user. ... It evaluates the worth of each candidate by comparing its predicted data with the actual input data. ... Thus, *test* selects the "best" explanation of the data.
>
> *Plan* produces direct (i.e., not chained) inference about likely substructures in the molecule from patterns in the data that are indicative of the presence of the substructure. In other words, *Plan* worked with combinatorially reduced abstracted sets to guide the search in a generally fruitful direction.

If some of the words in this description were replaced, the plan-generate-and-test approach seemed appropriate for the HASP tasks:

> *Generate* plausible ship candidates and their signal characteristics.
> *Test* by comparing the predicted signals with the real signals.
> *Plan* by selecting types of ships that could be in the region of interest. The Plan phase would use intelligence reports, shipping logs, and so on.

The system designers had already talked with the analysts and had read their training manuals. They knew that the necessary knowledge could be represented as rules, a form of domain knowledge representation that had proven its utility and power in DENDRAL. Difficulties were encountered immediately; some of these were:

1. The input data arrived in a continuous stream, as opposed to being batched like in DENDRAL. The problem of a continuous data stream was solved by processing data in time-framed batches.

2. The analysis of the activities in the ocean had to be tracked and updated over time.

Most importantly, past activities played an important role in the analysis of current activities.

3. There were numerous types of information that seemed relevant but remote from the interpretation process, for example, the average speeds of ships.

To address the second problem, it was immediately clear that a data structure was needed which was equivalent to a "situation board" used by the analysts; the data structure was called the Current Best Hypothesis (CBH). CBH reflected the most recent hypothesis about the situation at any given point in time. This could serve as the basis for generating a "plan;" that is, the CBH could be used as a basis for predicting the situation to be encountered in the next time frame. The prediction process could also utilize and integrate the variety of information mentioned in item 3. The predicted CBH would then be used (1) to verify that the interpretation from the previous time frame was correct, (2) to reduce the number of alternatives generated during past time-frames,[34] and (3) to reduce the number of new signals not accounted for in the predicted CBH that needed to be analyzed in full. CBH was thought of as a cognitive "flywheel" that maintained the continuous activities in a region of ocean between time frames. The initial design, a modified version of DENDRAL, was sketched out in December of 1973 (Figure 2-3).

t=t+1

| PLAN: generate preliminary classifications & cross correlation | HYPOTHESIZE: generate plausible platforms & sources | VALIDATE: compare simulated data and real data | PREDICT: generate probable situation in the next time frame |

| Preliminary hypothesis | Working hypothesis | Current best hypothesis | Predicted hypothesis |

(— Control flow; --- Data flow)
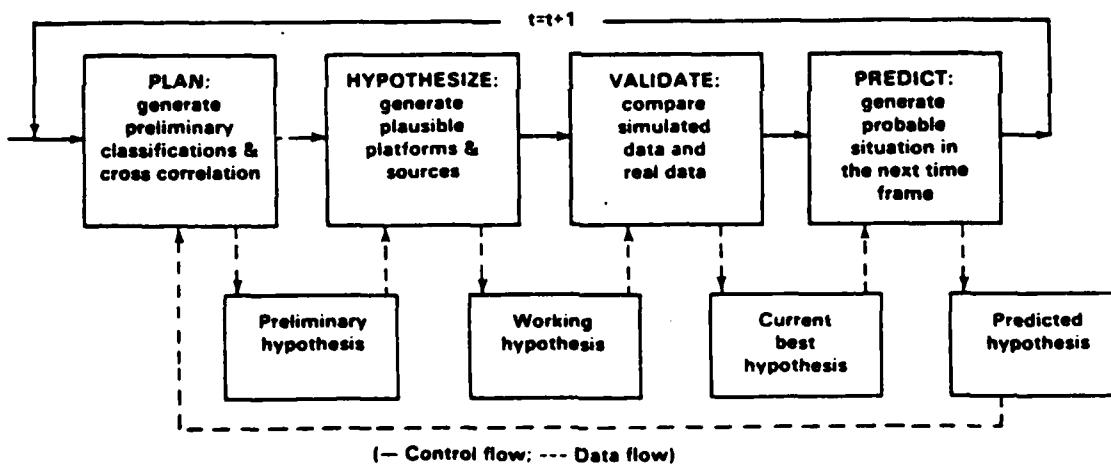
There was one global data structure that contained a hypothesis about the situation. After each of the plan, hypothesize, validate, and predict phases the hypothesis changed states. These states were called the preliminary hypothesis, the working hypothesis, the current best hypothesis, and the predicted hypothesis.

Figure 2-3:   HASP Design Based on DENDRAL

---

[34]There was only one solution hypothesis. However, some attributes, for example, platform type, could have alternative values.

Then there came the bad news: There was no plausible generator of the solution space, and there was no simulator to generate the signals of hypothesized platforms. The bad news had a common root; given a platform, there was a continuum of possible headings, speeds, and aspects relative to an array. Each parameter, in addition to variations in the water temperature, depth, and so on, uniquely affected the signals "heard" at an array. Consequently, there was a continuum of possibilities in the solution space as well as for the simulator to simulate. The designers tried to limit the number of possibilities, for example, by measuring the headings by unit degrees, but this left an enormous search space. Moreover, there was not enough knowledge to prune the space to make the generate-and-test method practical. The DENDRAL approach was abandoned. Then, the HEARSAY-II approach was learned of. The description of the approach produced enough of a mental shift in the way the HASP problem was viewed that a new solution could be designed. It should be noted in passing that HEARSAY-II in fact had generators and used them. It was the idea of fusing uncertain and partial solutions to construct solutions, combined with "island driving,"[35] that intrigued the designers.

The sonar analysts solved the problem piecemeal. They first identified a harmonic set in the signals. The "accounted-for" signals were then "subtracted" from the data set. Then another harmonic set would be formed with the remaining data and so on until all the signals were accounted for.[36] Each harmonic set implied a set of possible sources of sound (for example, a propeller shaft), which in turn implied a set of possible ship types from which the sounds could be emanating. Certain signal characteristics directly implied platform types, but this type of diversion from the incremental analysis was very rare. What the human analysts were doing was what might be called logical induction and synthesis.[37] Hypotheses were synthesized from pieces of data using a large amount of domain-specific knowledge that translated information in one form to information in another form, that is, transformed a description in one vocabulary to one in another vocabulary. For example, a set of frequencies was transformed

------

[35]*Island driving* is a problem solving strategy. A relatively reliable partial hypothesis is designated as an "island of certainty," and the hypothesis building pushes out from this solution island in many directions. This is sometimes called a "middle-out" strategy. There can be many islands of certainty driving the problem-solving process.

[36]As easy as it sounds, the task of harmonic set formation was a very difficult one, given noisy and missing data, and could produce large combinatorial possibilities. Addressing this problem became one of the major concerns in HASP.

[37]An interesting article on this point, *A More Rational View of Logic*, is by Alex P. Pentland and Martin A. Fischler; it appeared in the Winter 1983 issue of AI Magazine.

into a set of possible ship parts (for example, a shaft or a propeller) by using knowledge of the form, "If the harmonic set consists of ..., then it is most likely to be due to ..." . The partial solutions thus formed were then combined using other knowledge to construct acceptable solutions.

The analysts were also strongly model driven. There were common shipping lanes used by merchant ships traveling from one port to another. These platforms usually maintained a steady speed and heading. This and similar knowledge served to constrain the number of possible partial solutions. For example, if a hypothetical surface platform traveled across a shipping lane, then the possibility that it might be a merchant ship could be eliminated. In this example, a model of ship movements was able to aid in the platform classification process. Moreover, knowledge about the characteristics of platforms was used to combine lower-level, partial solutions. Suppose a platform type was hypothesized from an acoustic source, for example, a propeller shaft. Knowledge about the platform type (a model) was then used to look for other acoustic sources (for example, an engine) belonging to that platform. This type of top-down, model-driven analysis was used as often as the bottom-up signal analysis.

Once it was clear that interpretation in HASP, as in HEARSAY, was a process of piecemeal generation of partial solutions that were combined to form complete solutions, the HEARSAY-II system organization could be exploited. The CBH was partitioned into levels of analysis corresponding to the way analysts were used to thinking (that is, harmonic sets, sources, and ship types). The rule-based knowledge gathered for the purposes of pruning and guiding the search process was organized into sets of rules (knowledge sources) that transformed information on one level to information on another level.[38]

Nothing is as easy as it appears. There were many differences between the speech and the sonar signal understanding tasks that drove the HASP system architecture in a different direction from HEARSAY-II. The use of the blackboard as a situation board that evolves over time has already been mentioned. This is somewhat equivalent to asking a speaker to repeat his utterance over and over again while moving around, and having the interpretation improve with each repeated utterance as well as being able to locate the speaker after each utterance. After each utterance, the CBH would reflect the best that the system could do up to that point.

---

[38]It is interesting to note that many of the pieces of knowledge intended for pruning purposes could be converted into inductive knowledge. For example, a pruning rule that read "If a signal is coming from outside the normal traffic lane, then its source could not be cargo or cruise ships." could be used directly for reducing alternatives or could be converted to read "..., then its source is either military ships or fishing boats." One can hold the view that this is not surprising, because knowledge is knowledge and what counts is how and when it's used.

It was also mentioned that sets of rules were used as opposed to procedures in HEARSAY-II.
to represent knowledge sources. Rules were chosen because they were used in DENDRAL and
in MYCIN.[39] This choice of knowledge representation had a great influence in simplifying the
HASP scheduler. The following characteristics influenced the final design of HASP.

Events: The concept of events is inherent in the HASP problem. For example, a certain type
of frequency shift in the signal would be an event that implied the ship was changing its
speed. An appearance or disappearance of a signal would be an event that implied a new ship
was on the scene or a known ship was getting out of the range of the sensors, or it implied an
expected behavior of certain types of ships. This inherent task characteristic made it natural
for the HASP system to be an event-based system; that is, *an occurrence of a particular event
implied that new information was available for some* a priori *determined knowledge source to
pursue.* The goals of the task dictated what events were significant and what were not. This,
in turn, meant that the programmer (the knowledge engineer of today) could *a priori* decide
what changes in the blackboard, that is, events, were significant for solving the problem (as
opposed to the system noticing every change). Furthermore, the only time a knowledge source
needed to be activated was when some events occurred that it knew about. These task
characteristics, together with the use of a rule-based knowledge representation, helped redefine
and simplify the task of the scheduler in the sense that each piece of knowledge was more or
less self-selecting for any given event.[40]

Temporal events: In HEARSAY-II "time" meant the sequence in which the words appeared in a
spoken sentence. Based on the sequence of words, one could predict or verify the appearance
of another set of words later or earlier in the sequence. In HASP time had different
connotations. In one sense, time was similar to the separate utterance in the hypothetical
repetitive utterances problem mentioned earlier. There was information redundancy, as well as
new and different information (no two utterances sound exactly the same), as time went on.
Redundancy meant that the system was not pressed to account for every piece of data at each
time frame. It could wait to see if a clearer signal appeared later, for example. Also, time
meant that the situation at any time frame was a "natural" consequence of earlier situations,
and such information as trends and temporal patterns (both signal and symbolic) that occur
over time could be used. One of the most powerful uses of time in this sense was the
generation and use of expectations of future events.

---

[39]This is a good example of the cultural influence. No other representation was even considered.

[40]The relationship between events within the task and "events" in the system are discussed in Section 3.2.

Multiple input streams: Aside from the digitized data from many hydrophones, HASP had another kind of input -- reports. Reports contained information gathered from intelligence or normal shipping sources. These reports tended to use descriptions similar to those used on the ship level on the blackboard (CBH). Whereas the ordinary data came in at the bottom level for both HEARSAY and HASP, HASP had another input "port" at the highest level. Given the input at this level, the system generated the kinds of acoustic sources and acoustic signatures it expected in the future based on information in its taxonomic knowledge base. This type of *model-based expectation* was one of the methods used to "fuse" report data with signal data.

Explanation: The purpose of explanation is to understand what is going on in the system from the perspective of the user and the programmer. Because the needs of the users are different from those of the programmers, explanation can take on many forms. Explanation for the user was especially important in HASP, because there was no way to test the correctness of the answer. The only way to test the performance of the system was to get human analysts to agree that the system's situation hypotheses and reasoning were plausible. CBH, with its network of evidential support, served to justify the hypothesis elements and their hypothetical properties. It served to "explain" the relationships between the signal data and its various levels of interpretation. The explanation of the reasoning, that is, "explaining" which pieces of knowledge had been applied under what circumstances, was made possible by "playing back" the executed rules.[41]

There were many other differences, but these characteristics had the most impact on the design of the eventual system. The list serves to illustrate how strongly the task characteristics influence blackboard architectures. (The details of the organization of the HASP system are explained in Section 3.2.)

Since Hearsay-II and HASP, there have been a variety of other programs whose system designs are rooted in the blackboard model. These programs include applications in the area of interpretation of electron density maps of protein crystals [50], planning [20], scene analysis [30], and signal understanding and situation assessment [48], [54]. There are other applications currently being built in the areas of process control, very large-scale intergration

---

[41]In MYCIN and other similar rule-based programs, explanation consists of a playback of rule firings. In HASP the ordinary method of playback turned out to be useful only to programmers for debugging purposes. For the user, the rules were either too detailed or were applied in a sequence (breadth first) that was hard for the user to understand. In HASP the explanation of the line of reasoning was generated from an execution history with the help of "explanation templates" that selected the appropriate rule activities in some easy-to-understand order.

(VLSI) design, crisis management, image understanding, and signal interpretation. Many applications in Defence Advacned Research Project Agency's (DARPA) Strategic Computing Program in the areas of military battle management, a pilot's associate, and autonomous vehicles utilize the blackboard model. To date, there is no commonly agreed upon architecture for blackboard systems. Rather, there are more or less strict interpretations of the model. The blackboard framework that distills the common constructs and features of many blackboard systems has been introduced. In the next section, documented systems that follow the intent and spirit of the blackboard model are described.

3. Blackboard Application Systems

The application systems described hare are presented in chronological order. The design of many of the systems is similar because of similarities in the application tasks, propagation of ideas, or involvement of the same designers. Figure 3-1 shows a general chronology and intellectual lineage of the various application and skeletal systems. The figure includes some of the better-known and better-documented systems. Only a few of the many application systems are described here; they were chosen because they illustrate different designs and because they contributed new ideas and features to the repertoire of blackboard system architectures. For each application, the task and domain characteristics are described. The description is followed by a summary of the system design in four parts: the blackboard structure, the knowledge source organization, the control component, and the knowledge application strategy employed. Unique features in the system are pointed out and discussed within the context of either the application task or its history.
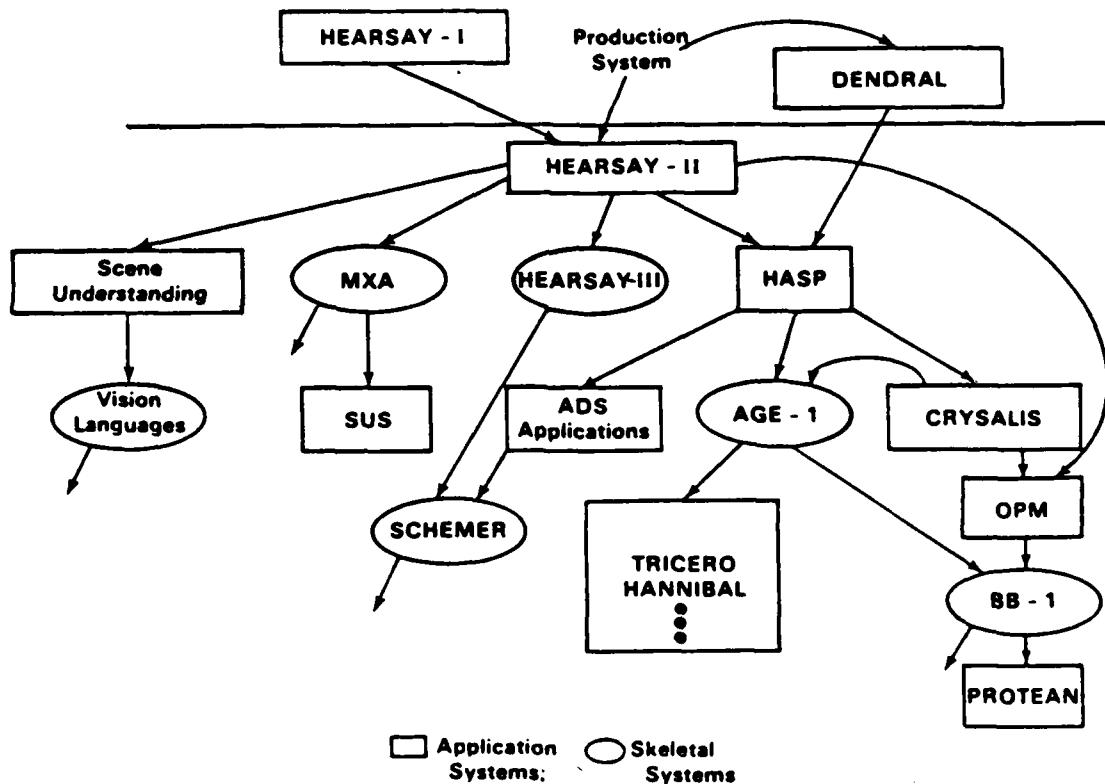
References: *Hearsay-1 [39]; Production system [32] and [6]; Dendral [27]; Hearsay-II [8]; Scene understanding [30]; Vision language [43]; MXA [23]; SUS [23]; HEARSAY-III [9]; HASP [36]; ADS (Advanced Decision System, Inc.) applications [48] and [29]; AGE-1 [35]; CRYSALIS [50]; OPM [20]; BB-1 [18]; TRICERO [53]; HANNIBAL [3]; PROTEAN [19].*

Figure 3-1:   Influences Among Blackboard Systems

3.1. HEARSAY-II


"IS THE SYSTEM RUNNING?"
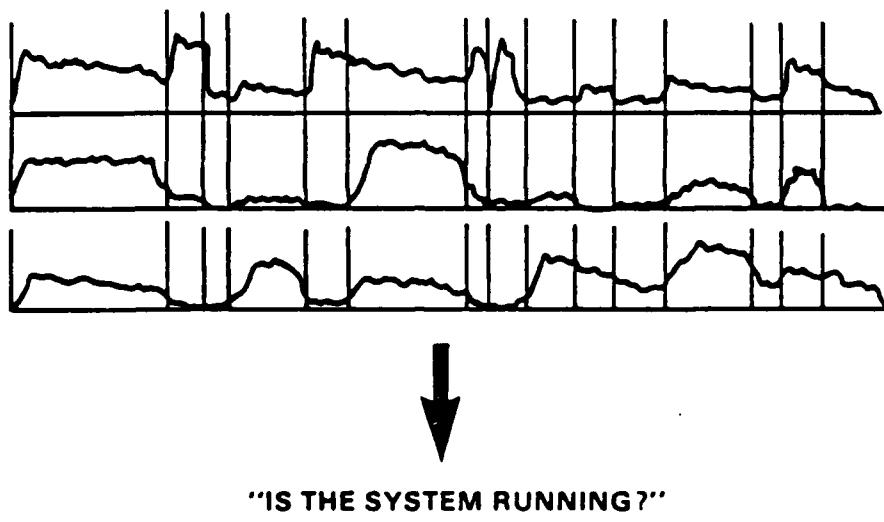
Figure 3-2: The HEARSAY-II Task

Most of the background information on HEARSAY-II was covered in Section 2.2 and is not be repeated here. One additional item of historical context is worth noting, however. Various continuous speech-understanding projects were brought under one umbrella in the Defense Advanced Research Projects Agency (DARPA) Speech Understanding Project, a five-year project that began in 1971. The goals of the Speech Understanding Project were to design and implement systems that "accept continuous speech from many cooperative speakers of the general American dialect in a quiet room over a good quality microphone, allowing a slight tuning of the system per speaker, by requiring only natural adaptation by the user, permitting a slightly selected vocabulary of 1,000 words, with a highly artificial syntax...in a few times real time..." [33] Hearsay-II was developed at Carnegie-Mellon University for the Speech Understanding Project and successfully met most of these goals.

The Task

The goal of the HEARSAY-II system was to understand speech utterances. To prove that it understood a sentence, it performed the spoken commands. In the earlier HEARSAY-I period, the domain of discourse was chess (for example, bishop moves to king knight five). In the HEARSAY-II era, the task was to answer queries about, and to retrieve documents from, a collection of computer science abstracts in the area of artificial intelligence. For example, the system understood the following types of command:

"Which abstracts refer to the theory of computation?"

"List those articles."
"What has McCarthy written since nineteen seventy-four?"

The HEARSAY-II system was not restricted to any particular task domain. "Given the syntax and the vocabulary of a language and the semantics of the task, it attempts recognition of the utterance in that language." [38] The vocabulary for the document retrieval task consisted of 1011 words in which each extended form of a root, for example, the plural of a noun, was counted separately. The grammar defining a legal sentence was context-free and included recursion, and imbedded semantics and pragmatic constraints. For example, in the place of noun in conventional grammars, this grammar included such non-terminals as topic, author, year, and publisher. The grammar allowed each word to be followed, on the average, by seventeen other words in the vocabulary.

The problem of speech understanding is characterized by error and variability in both the input and the knowledge. "The first source of error is due to deviation between ideal and spoken messages due to inexact production [input], and the second source of error is due to imprecise rules of comprehension [knowledge]." Because of these uncertainties, a direct mapping between the speech signals and a sequence of words making up the uttered sentence is not possible. The HEARSAY designers structured the understanding problem as a search in a space consisting of complete and partial interpretations. These interpretations were organized within an abstraction hierarchy containing signal parameters, segments, phones, phonemes, syllables, words, phrases, and sentence levels. This approach required the use of a diverse set of knowledge that produced large numbers of partial solutions on the many levels. Furthermore, the uncertainties in the knowledge generated many competing, alternative hypothetical interpretations. To avoid a combinatorial explosion, the knowledge sources had to construct partial interpretations by applying constraints at each level of abstraction. For example, one kind of constraint is imposed when an adjacent word is predicted, and the prediction is used to limit subsequent search. The constraints also had to be added in such a way that their accrual reduced the uncertainty inherent in the data and the knowledge sources.

In order to control the combinatorial explosion and to meet the requirement for near real-time understanding, the interpretation process had to be selective in exploiting the most promising hypotheses, both in terms of combining them (for example, combining syllables into words) and in terms of predicting neighboring hypotheses around them (for example, a possible adjective to precede a noun). Thus, the need for incremental problem solving and flexible, opportunistic control were inherent in HEARSAY's task.

The Blackboard Structure

The blackboard was partitioned into six to eight (depending on the configuration) levels of analysis corresponding to the intermediate levels of the decoding process.[42] These levels formed a hierarchy in which the solution-space elements on each level could be described loosely as forming an abstraction of information on its adjacent lower level. One such hierarchy was comprised of, from the lowest to the highest level: parametric, segmental, phonetic, phonemic, syllabic, lexical, phrasal, and conceptual levels (see Figure 3-3). A blackboard element represented a hypothesis. An element at the lexical level, for example, represented a hypothesized word whose validity was supported by a group of syllables on the syllable level. The blackboard could be viewed as a three-dimensional problem space with time (utterance sequence) on the x-axis, information levels containing a hypothesized solution on the y-axis, and alternative solutions on the z-axis. [24]

Each hypothesis, no matter which level it belonged to, was constructed using a uniform structure of attribute-value pairs. Some attributes, such as its level name, were required for all levels. The attributes included a validity rating and an estimate of the "truth" of the hypothesis represented as some integer value. The relationships between the hypotheses on different levels were represented by links, forming an AND/OR tree over the entire hierarchy. Alternative solutions were formed by expanding along the OR paths. Because of the uncertainty of the knowledge sources that generated the hypotheses, the blackboard had a potential for containing a large number of alternative hypotheses.

The Knowledge Source Structure

Each knowledge source had two major components: a condition part (often referred to as a precondition) and an action part. Both the condition and the action parts were written as arbitrary SAIL procedures. "The condition component prescribed the situations in which the knowledge sources may contribute to the problem-solving activity, and the action component specified what that contribution was and how to integrate it into the current situation." [8] When executed, the condition part searched the blackboard for hypotheses that were of interest to its corresponding action part; all the relevant hypotheses found during the search were passed on to the action part. Upon activation, the action part processed all the hypotheses passed to it. The tasks of the knowledge sources ranged from classification (classifying acoustic segments into phonetic classes), to recognition (recognizing words) to generation and

---

[42]See [25] for a comprehensive discussion on the results of experiments conducted with two different blackboard configurations.

LEVELS   Knowledge Sources

PHRASAL
                                                    Syntatic-Semantic Parser
                                                    Syntatic-Semantic Hypothesizer
LEXICAL
                                                    Phoneme Hypothesizer

SYLLABIC
                                                    Word Candidate Generator

SURFACE-
PHONEMIC
                                                    Phone-Phoneme Synchronizer

PHONETIC
                                                    Phone Synthesizer

                                                    Segment Combiner
SEGMENTAL
                                                    Segmenter-Classifier

PARAMETRIC
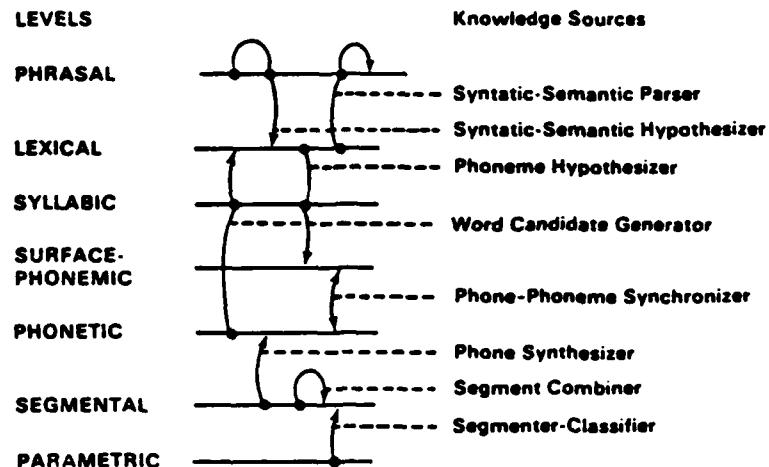
Figure 3-3: HEARSAY-II Blackboard and Knowledge Sources

evaluation of predictions.

Control

The control component consisted of a blackboard monitor and a scheduler (see Figure 3-4). The monitor kept an account of each change made to the blackboard, its primitive change type, and any new hypotheses. Based on the change types and declarative information provided by the condition part of the knowledge sources, the monitor placed pointers to those condition parts which potentially could be executed on a scheduling queue.[43] In addition to the condition parts ready for execution, the scheduling queue held a list of pointers to any action parts ready for execution. These actions parts were called the *invoked knowledge sources*. A knowledge source became invoked when its condition part was satisfied. The condition parts and the invoked knowledge sources on the scheduling queue were called *activities*. The scheduler calculated a priority for each activity at the start of each system cycle and executed the activity

---

[43]In figure 3-4 the "Focus-of-control database" contained a table of primitive change types and the condition parts that could process each change type. The primitive change types possible within the system were predefined and consisted of such items as "new syllable" and "new word created bottom up." This paragraph is based on discussions with Lee Erman.

with the highest priority in that cycle.

In order to select the most productive activity (the most important and promising with the least amount of processing and memory requirements), the scheduler used experimentally derived heuristics to calculate the priority. These heuristics were represented as imbedded procedures within the scheduler. The information needed by the scheduler was provided in part by the condition part of each invoked knowledge source. The condition part provided a *stimulus frame*, a set of hypotheses that satisfied the condition; and a *response frame*, a stylized description of the blackboard changes the knowledge source action part might produce upon execution. For example, the stimulus frame might indicate a specific set of syllables, and the response frame would indicate an action that would produce a word. The scheduler used the stimulus-response frames and other information on the blackboard to select the next thing to do.
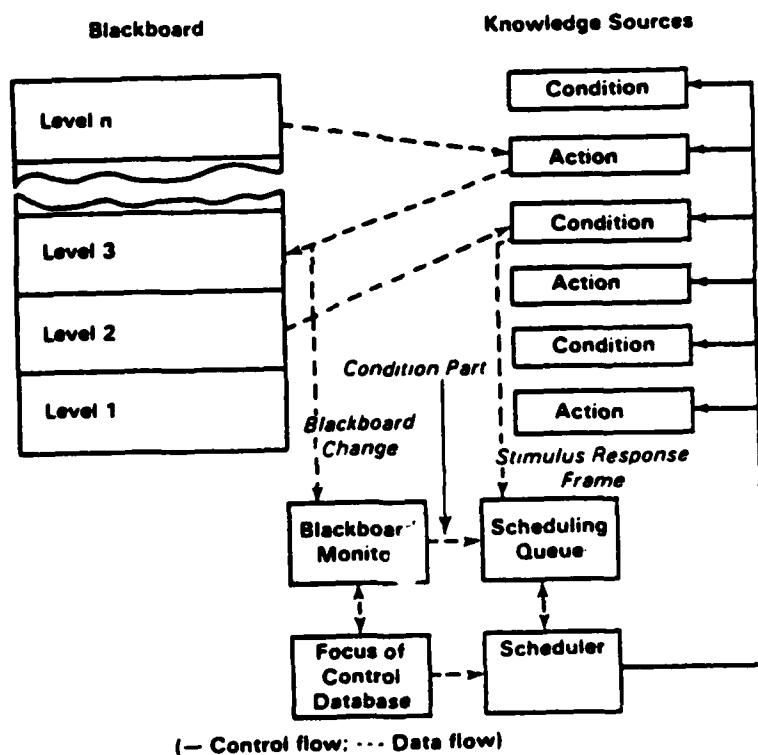

Figure 3-4:   Schematic of HEARSAY-II Architecture

The control component iteratively executed the following basic steps:

1. The scheduler selected from the scheduling queue an activity to be executed.

2. If a condition part was selected and executed and if it was satisfied, a set of

> stimulus-response frames was put on the scheduling queue together with a pointer
> to the invoked knowledge source.

3. If an action part was selected and executed, the blackboard was modified. The
   blackboard monitor posted pointers to the condition parts that could follow up the
   change on the scheduling queue .

The problem of *focus of attention* was defined in the context of this architecture as a problem
of developing a method which minimized the total number of knowledge source executions and
which achieved a relatively low rate of error. The focus-of-attention problem was viewed as a
knowledge-scheduling problem as well as a resource-allocation problem.[44] In order to control
the problem-solving behavior of the system, the scheduler needed to know the goals of the task
and the strategies for knowledge application to be able to evaluate the next best move.
Although various general solutions to this problem have been suggested [15], it appears that
ultimately one needs a knowledge-based scheduler for the effective utilization of the knowledge
sources.[45]

Knowledge-Application Strategy

Within the system framework described earlier, HEARSAY-II employed two problem-solving
strategies. The first was a *bottom-up* strategy whereby interpretations were synthesized directly
from the data, working up the abstraction hierarchy. For example, a word hypothesis was
synthesized from a sequence of phones. The second was a *top-down* strategy in which
alternative sentences were produced from a sentential concept, alternative sequences of words
from each sentence, alternative sequences of phones from each word, and so on. The goal of
this recursive generation process was to produce a sequence on the parametric level that was
consistent with the input data (that is, to generate a hypothetical solution and to test it against
the data). Both approaches have the potential for generating a vast number of alternative
hypotheses and with it a combinatorially explosive number of knowledge source activations.
Problem-solving activity was, therefore, constrained by selecting only a limited subset of
invoked knowledge sources for execution. The scheduling module thus played a crucial role
within the HEARSAY-II system.

---

[44]If we compare the HEARSAY-II control constructs with those of the blackboard framework discussed in Section
1.2, they are basically the same. Some aspects of the control in HEARSAY are emphasized more (for example,
scheduling) than others.

[45]The current work of Barbara Hayes-Roth on the BB-1 system elaborates this point [18].

Orthogonal to the top-down and bottom-up approaches, HEARSAY-II employed a general
hypothesize-and-test strategy. A knowledge source would generate hypotheses, and their validity
would be evaluated by some other knowledge source. The hypothesis could be generated by a
top-down analytic or a bottom-up synthetic approach. Often, a knowledge source generated or
tested hypotheses by matching its input data against a "matching prototype" in its knowledge
base. For example, a sequence of hypothesized phones on the phone level were matched
against a table containing prototypical patterns of phones for each word in the vocabulary. A
word whose phones satisfied a matching criterion became a word hypothesis for the phones.
The validation process involved assigning credibility to the hypothesis based on the consistency
of interpretation with the hypotheses on an adjacent level.

At each problem-solving step, any one of the bottom-up synthesis, top-down goal generation,
neighborhood prediction, hypothesis generation, and hypothesis evaluation might have been
initiated. The decision about whether a knowledge source could contribute to a solution was
local to the knowledge source (precondition). The decision about which knowledge source
should be executed in which one of many contexts was global to the solution state (the
blackboard), and the decision was made by a global scheduler. The scheduler was opportunistic
in choosing the next step, and the solution was created one step at a time.

Additional Notes

1. The condition parts of the knowledge sources were complex, CPU-intensive procedures that
needed to search large areas of the blackboard. Each knowledge source needed to determine
what changes had been made since the last time it viewed the blackboard. To keep from firing
the condition parts continually, each condition part declared *a priori* the kinds of blackboard
changes it was interested in. The condition part, when executed, looked only at the relevant
changes since the last cycle. All the changes that could be processed by the action part were
passed to it to avoid repetitive executions of the action part.

2. The HEARSAY-II system maintained alternative hypotheses. However, the maintenance and
the processing of alternatives are always complex and expensive, especially when the system
does not provide a general support for this. In HEARSAY-II, the problem was aggravated by
an inadequate network structure that did not allow the shared network to be viewed from
different perspectives. In the current jargon, it did not have good mechanisms for processing
multiple worlds.[46]

---

[46]Currently, there are better techniques for processing and maintaining alternative worlds. However, these techniques
have yet to be integrated into blackboard systems.

3. The evidence to support a hypothesis at a given level can be found on lower levels or on higher levels. For example, given a word hypothesis, its validity could be supported by a sequence of syllables or by grammatical constraints. The evidential support is represented by directional links from the evidence to the hypothesis it supports. The link that goes from a higher-level to a lower-level hypothesis represents a "support from above" (that is, the justification for the hypothesis can be found at a higher level). A link that goes in the opposite direction represents support from below (that is, the reason for the hypothesis can be found at a lower level). Although the names of the support mechanisms were first coined in HASP [34], the bidirectional reasoning mechanisms were first used in the HEARSAY-II system.

4. In HEARSAY-II the confidence in a hypothesis generated by a knowledge source was represented by an integer between 1 and 100. The overall confidence in the hypothesis was accumulated by simple addition of the confidence attached to the evidence (that is, supporting hypotheses). When the confidence in a hypothesis was changed, the change was propagated up (if the support was from below) and down (if the support was from above) the entire structure.

3.2. HASP/SIAP
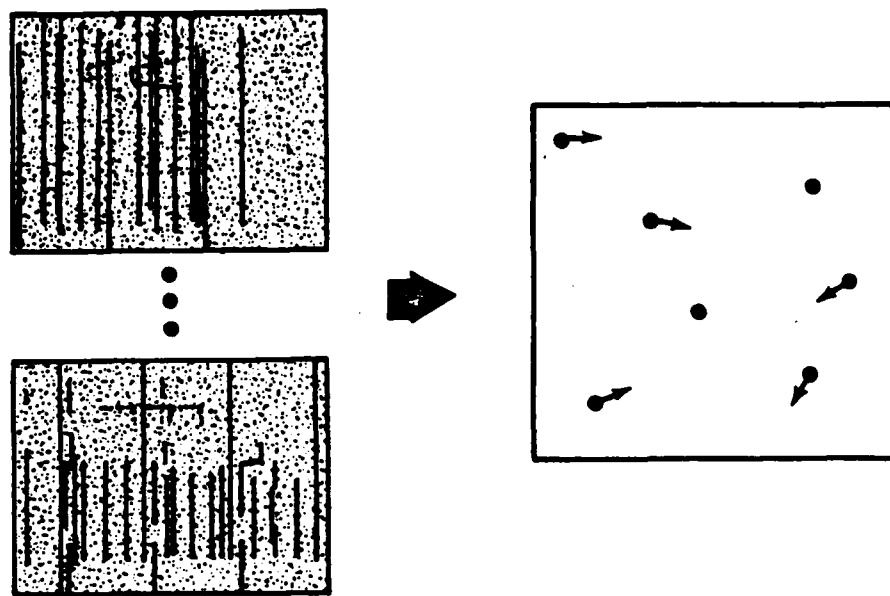
Figure 3-5:   HASP/SIAP Task

The HASP project began in 1972 under the sponsorship of the DARPA. The HASP project was terminated in 1975 but was reinstated in 1976 under the name SIAP. At the time, the computational resources needed to maintain a major ocean surveillance system of sensors with conventional methods of statistical signal processing seemed economically unfeasible. It was also a time when artificial intelligence techniques were first being applied to the problem of signal interpretation. The DENDRAL program [27] was achieving significant success and the Speech Uriderstanding Project, of which the HEARSAY Project was a part, was under way. The major objectives of the HASP project were to demonstrate that artificial intelligence techniques could contribute significantly in addressing the surveillance problem and, further, that the task could be accomplished with reasonable computing resources. HASP was successful in meeting both these objectives.

The Task

The task of the HASP/SIAP system was to develop and maintain a situation board that reflected the activities of platforms (surface ships and submarines) in a region under surveillance. The situation board was developed by interpreting multiple, continuous streams of acoustic signals produced by objects in the region and by integrating intelligence reports with the interpretation.

The acoustic input to the system came in the form of digitized data from multiple hydrophone

arrays, each monitoring a part of the region.[47] Each array had multiple hydrophones with some directional resolution. The major sources of acoustic radiation were rotating shafts and propellers and reciprocating machinery on board a platform. The *signature*, or sound spectrum, of a platform under steady operation contained persistent fundamental narrow-band frequencies and certain of their harmonics. The front-end signal-processing hardware and software detected energy peaks appearing at various spectral frequencies and followed these peaks over time. On an analyst's sonograms, the peaks appeared as a collection of dark vertical stripes (see Figure 3-5). Under ideal conditions, a hydrophone picked up sound energy near its axis. In practice, the terrain of the ocean floor, water temperature, and other platforms interfered, producing signals with very low signal-to-noise ratios. That is, the stripes in the sonogram appeared against a very fuzzy background.

In addition to the acoustic data, intelligence reports were available to HASP. The reports contained information about movements of friendly and hostile platforms with varying degrees of confidence. Routine information on commercial shipping activities was also included in these reports.

As in the speech understanding problem, the sonar signal-understanding problem is characterized by a large solution space, a low signal-to-noise ratio, and uncertain knowledge. Unlike the speech problem, the semantics and the syntax are ill defined in the sonar problem. That is, the targets of highest priority, the enemy submarines, are most likely to be ill understood and, at the same time, are trying their best to go undetected. The implications are these: (1) There is no "legal move generator" for the solution space except at the highest level of abstraction. (It is assumed that different types of enemy submarines and their general characteristics are known.) (2) One must rely heavily on the analysts' methods and heuristics in detecting and classifying enemy submarines. (3) In order to find the targets, the analysts accounted for all known entities (primarily surface platforms) and looked for the targets of interest within the unaccounted-for data.[48] The problem is somewhat akin to the following tasks: When there are two people talking at the same time, one in English and another in a relatively unfamiliar language, try to pick up what the non-English speaker is saying. Another

---

[47]During the first phase of the project, the acoustic input consisted of segments that described signal events. For example, a piece of input might have contained a frequency and indicated it as a beginning of a frequency shift (called a *knee*). Later, five-minute segments produced by a signal-processing front-end system were used as the input data.

[48]This does not guarantee that the targets will be found. For one thing, the targets might be very quiet, and their sound might not be picked up by the hydrophones, or their sound might be overshadowed by noisier platforms. The HASP system conjectured about their existence and their whereabouts from other information.

task is the cocktail conversation problem in which many people are talking as they move around; the task is to keep track of each person using data from microphones scattered around the room.

Even given all the difficulties, there were aspects of the problem that made it tractable. The situation unfolded over a relatively long period of time because the platforms moved rather slowly, but the data collection was relatively frequent and from many different locations. This meant that the system was given many chances to interpret the situation with data sets containing slightly different information. For example, two hydrophones might pick up incomplete harmonic sets attributable to the same platform, but they might be fractured in different ways. When combined, they provided more information than from each one separately. There also were many different kinds of knowledge that could be used, bits and pieces, such as in the Koala problem discussed in Section 1.1. The general strategy employed was to accumulate both positive and negative evidence for a hypothesis element.

The Blackboard Structure

The data structure on the blackboard represented the best understanding of the situation at any given point in time. It was a dynamic entity that evolved over time. Referred to as the *current best hypothesis* (CBH), it was partitioned into an abstraction hierarchy consisting of input segments, lines, harmonic sets, acoustic sources, platforms, and fleet levels (see Figure 3-6). The signal data arrived on the segments level, and the report data arrived on either the fleets level or the platforms level, depending on the content of the report.

Unlike the HEARSAY-II system in which the "answer" to the problem was the hypothesized sentence on the highest level, HASP's "answer" was the network of partial solutions that spanned the entire blackboard. In other words, partial solutions were considered acceptable, if not desirable, solutions. For example, a partial solution of the form, "There's something out there producing these lines," was acceptable, even though a preferable solution was, "There is a platform of type x, whose engine is accounted for by the following harmonics and whose propeller seems to be producing the following lines, and no shaft data are currently being received."

The nodes on the blackboard were called *hypothesis elements* rather than "hypotheses" as they were in HEARSAY-II. The hypothesis elements formed a network, each element representing a meaningful aggregation of lower-level hypothesis elements. No attempt was made to maintain uniformity of attributes across the levels. Each knowledge source knew the relevant vocabulary (*attributes*) associated with those levels in which it was interested. The lines level and the harmonic-sets level used a descriptive vocabulary that dealt primarily with signal

characteristics, and the sources level used vocabulary dealing primarily with machinery. Thus, the point of signal-to-symbol transformation can be said to have occurred between the harmonic-sets level and the sources level. Signal information in a hypothesis element on the harmonic sets level was translated into machinery information in a hypothesis element on the sources level, that is, there was an element-for-element translation between the two levels.

In contrast to HEARSAY-II, each hypothesis element could have alternative values for its attributes but no alternative links. The hierarchy was organized as an AND tree, with a possibility for local alternatives. Although this approach reduced computational time and space, it was awkward for the system to "change its mind" about the solution. In HEARSAY, changing its mind might only have involved focusing on an alternative structure. In HASP either the affected hypothesis elements had to be reanalyzed (which could result in reorganizing the whole CBH), or the past analyses dealing with the elements in question had to be forgotten and the analysis restarted from the point of departure. The latter approach was used in HASP because the human analysts tended to behave in a similar manner.[49]
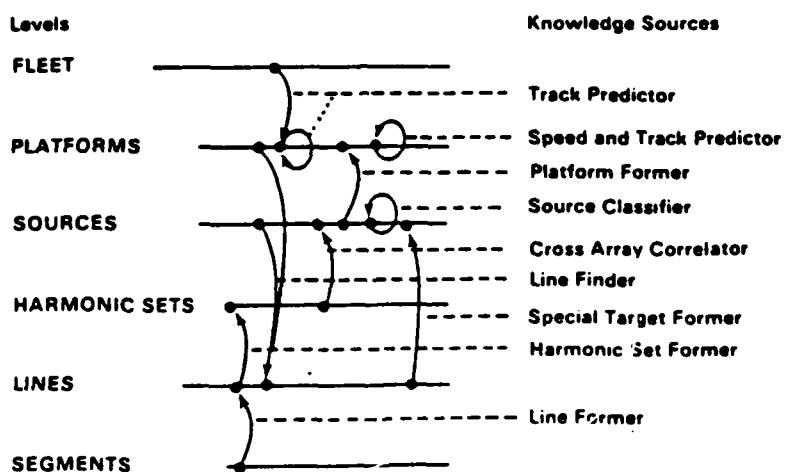
Figure 3-6: HASP/SIAP Blackboard and Knowledge Sources

In addition to the blackboard, HASP had other globally accessible information generated directly or indirectly by the knowledge sources (refer to Figure 3-7). This global information

---

[49]In the human system there are analysts whose task is to do off line postanalyses. What they learn from the postanalyses is often added to the pool of knowledge about the task. HASP had no counterpart to this activity.

was used primarily by the control modules:

Event List: All changes made to the blackboard, together with the types of these changes, were posted on the event list. Each event has a generic "change type" associated with it. An event also had associated with it a particular blackboard node (hypothesis element). An event in the event list was selected by a control module to become a focus of attention. The focus of attention then had two implicit components: a change type and a blackboard node. (A more detailed discussion can be found in the Control section.)

Expectation List: The expectation list contained events (event types and associated hypothesis elements) that were expected to occur in the future. Thus, acoustic signature of platforms reported to be in the region in the intelligence reports were posted on the expectation list. The canonical acoustic signatures of all the known platforms were stored in a static knowledge base.[50] Periodically the expectation list was searched to see if expected data had arrived.

Problem List: This list contained a description of the various problems the knowledge sources encountered. For example, when no rule fired during the execution of a knowledge source, it might have meant, "I should know, but don't." Such information was useful to the programmers. The most important use of this list, however, was for posting missing or desired information. A knowledge source could post pieces of information, that if available, would increase the confidence in its hypothesis. For example, a knowledge source might indicate that if the dependency relationship was known for a given set of lines, it might be able to identify the platform. In such a case, an operator might provide the information if it was known, or a goal might be set up by a control module to find the information.

Clock-Event List: A clock event consisted of a time and associated rules. The rules were to be executed at the designated time. Because behaviors at various levels were known for some types of platforms, knowledge sources tracked the expected and actual behavior by this mechanism. The types of behavior known to the system ranged from the temporal characteristics of the sonograms to the physical movement of the platforms.

History List: All the processed events and their context (for example, a blackboard node and its values and the bindings in a rule that made the change) were kept on this list. The history list was used to recount the knowledge-application steps that

---

[50]In the entire discussion of blackboard systems, the role and the form of the static knowledge base have been omitted. It is assumed that taxonomies, facts, and definitions are represented in some form. This type of knowledge is awkward to represent as rules and is usually represented as tables, records, property lists, or frames.

led to the generation of the CBH.[51] This list was also used by the programmers to ensure that solutions were arrived at by an expected line of reasoning. We wanted to detect occurrences of right answers for wrong reasons.
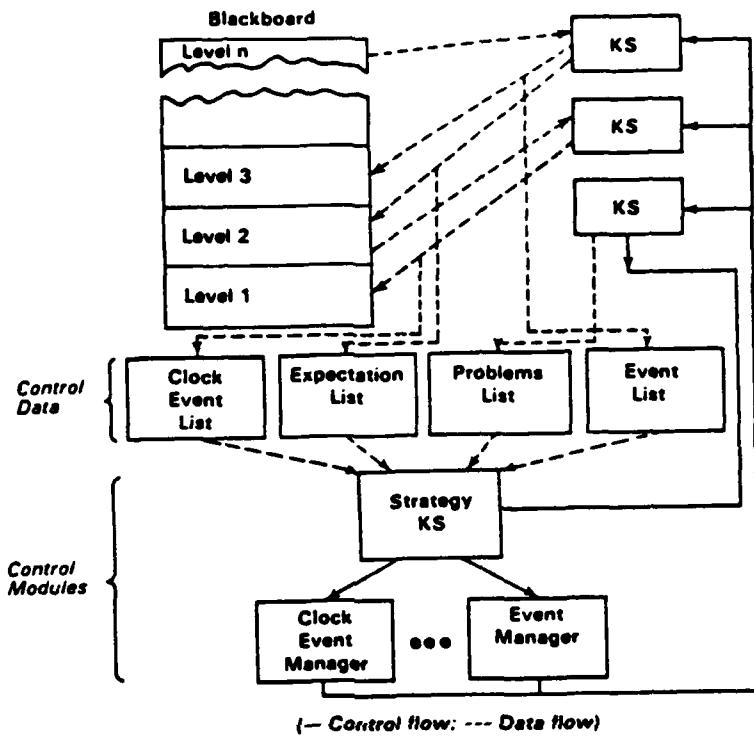

Figure 3-7:   HASP/SIAP System Organization

The Knowledge Source Structure

Each knowledge source consisted of a precondition part and an action part. In contrast to the HEARSAY-II knowledge source organization, the precondition part and the action part were contained in one module. The precondition part consisted of a list of pairs of tokens; the pair consisted of a name of an event type and its modifier (new, old, or modified). The modifier indicated the status of the hypothesis element (for example, modified hypothesis element) that was the focus of the events. When an event became "focused," knowledge sources whose precondition contained the event type of the focused event were executed. An event type was one of several predefined category of changes that could be made in the system. The action part consisted of a set of rules. In this knowledge source organization, the precondition can be

---

[51]Because the program processed events in a breadth-first order, and humans had difficulties in following this processing order, the history list was used to construct a text that made it appear as though the processing had been in depth-first order.

viewed as a simple trigger for a set of rules. The detailed test for applicability of knowledge were in the condition parts of the rules. The knowledge source could create bindings local to it that remained valid for the duration of its execution. The bindings served to "freeze" the context until all the rules in a knowledge source were evaluated.

Control

Each of the control modules in HASP was written in the same form as the domain knowledge sources, that is, as a set of rules. The knowledge sources formed a simple control hierarchy (see the control modules in Figure 3-7). Although the control knowledge sources were logically independent, they were executed in a predefined order. The strategy knowledge source decided which categories of events (that is, clock-event, problem, expectation, or blackboard) to process next, based on priorities as encoded in its knowledge base. An appropriate event-management knowledge source was executed based on this decision. Once activated, an event manager, in turn, decided which specific event to focus on. The basis of this decision varied with the event manager. For example, the clock-event manager selected events that needed to be processed at a given time, and within those events the priority rested with events dealing with enemy platforms. The knowledge sources associated with the focused event were then executed. The node associated with the focused event served as the context for the knowledge source's execution. For example, the strategy knowledge source might have decided that it was time to process a blackboard event. It would activate the blackboard-event manager. The event manager in turn looked through the event list and selected an appropriate event as the focus of attention. Finally, one or more knowledge sources whose precondition contained the blackboard change type of the focused event were executed. The node associated with the event (that is, the blackboard node containing a change) served as the context for the knowledge sources.

It was mentioned in Section 2.3 that the scheduling module and the focus-of-attention mechanisms were simpler in HASP than in HEARSAY-II. This is true in view of the following: Because it was known what blackboard changes were significant for making progress toward a solution, the HASP programmer decided what blackboard changes were to be called "events." That is, only certain changes to the blackboard were called events. For each such change, it was also known what knowledge sources were available for following up on the new information. By making the precondition of a knowledge source the occurrence of specific types of blackboard changes, the selection of knowledge sources for a given event became a very simple matter. For example, a table of change types and applicable knowledge sources

could be used.[52] In this scheme, however, the process of selecting the most promising event (a blackboard change and the node on which the change was made) became a major issue. The selection of an event is really the selection of a node, which, in turn, is really a selection of a solution island. Thus, the focus-of-attention problem in HASP was primarily a problem of determining which solution island to work on next, rather than a problem of which knowledge source to apply next, as in HEARSAY-II. In HASP the hierarchical control knowledge sources were all biased toward the selection of a solution island to be pursued that would have the highest payoff in subsequent processing cycles. Once the focus-of-attention event was selected, the relevant knowledge sources were easily selected, and all the knowledge sources were executed in a predetermined but interchangeable order.

The basic actions of the control component were iterations of the following:

1. The strategy knowledge source decided which event category to focus on, that is, clock events, expectations, problems, or blackboard events.

2. The manager of the chosen event category selected a specific event from that category to process next. The event information contained the name of a node to which a change was made and a change type associated with that change. The node name and the change type constituted the focus of attention.

3. Based on the change type of the focus of attention node, knowledge sources associated with the change type were executed. The node associated with the focus of attention served as the context for the activation of the knowledge sources.

4. The executing knowledge sources produced changes to the blackboard and the changes were recorded.

To summarize, in HASP there were four categories of events: expectation, clock, problems and blackboard. Each category of events contained a predetermined set of event types (that is, a set of expectation event types, a set of blackboard-change event types, and so on). For each event type, the knowledge sources that could process an instance of the event type were also predetermined. The system was openended in that new event types and new knowledge sources could be added without perturbing the existing ones. The major task for the control mechanism was to select the next solution island to be investigated.

---

[52]In HASP, a set of simple rules was used. The condition side contained event types and a few other simple conditions, and the action side contained a sequence of knowledge sources to be executed.

Knowledge-Application Strategy

As in HEARSAY-II, HASP used several problem-solving approaches. A basic generate-and-test method was used to generate hypothesis elements and to test their credibility. Instead of using a *legal move generator* as was the case in HEARSAY-II where the space of legal solutions was known from the grammar and vocabulary, HASP used a *plausible move generator* based on the heuristics used by the analysts. The construction of higher-level partial solutions from lower-level partial solutions, the determination of their properties, the generation of expectations, and so on, were driven by empirical association rules obtained from an analyst.

Most of the forty to fifty knowledge sources in HASP were engaged in bottom-up processing. Several pieces of data from a lower level were combined to form or update information on a higher level (for example, lines into harmonic sets). Similarly, information on one level was translated into a different vocabulary on another level (for example, harmonic sets into mechanical parts). The data were processed breadth first. That is, all the harmonic sets were formed from lines, and all sources were assigned to harmonic sets, and so on, in a pipeline fashion up the hierarchy.

The most powerful reasoning strategy used in HASP was the top-down, model-driven strategy. The assumption underlying model-driven reasoning is the following: In the interpretation of data, the amount of processing can be reduced by carefully matching selected pieces of data with discriminating or important features of a model (a frame or a script). A successful match tends to confirm the model as an explanatory hypothesis for the data. In a continuous-data interpretation task, the model, combined with periodic confirmatory matches, serves as the "cognitive flywheel" that maintains the ongoing "understanding." In driving a car, for example, our model of the road situation (prototypical highway characteristics, shapes of cars, their range of speed, their normal behavior, and so forth) saves us from continually having to process every bit of data within our visual range. Therefore, we don't "notice" the color of the upholstery of the car in front of us even though that piece of information is often available. The danger with this approach is that data can often match a wrong model for a long time, especially when the discriminating features are not carefully chosen.[53]

---

[53]How often have you listened to a person and thought that person was talking about a particular topic before suddenly realizing it was a different topic all the time? (See [1] for a simple experiment relevant to this topic.) The same pieces of knowledge from the PUFF [22] program were used in data-driven, goal-driven, and model-driven approaches. Although the model-driven approach ran the fastest, extra knowledge had to be added to keep it from making the wrong diagnoses.

With this caveat, a model-driven approach is a very powerful device in interpreting noisy data. In HASP a model-driven approach was used quite extensively and successfully. For example, it was used in determining which lines formed a harmonic set. In fact, the CBH served as a situation model from one time frame to the next. There was an implicit assumption that the current state of affairs was not significantly different from the state a few minutes earlier. To make this assumption work, HASP focused on finding counterevidence for a hypothesis as much as on finding supporting evidence.

Within an abstraction hierarchy, model-driven reasoning is usually a top-down process. For example, if a platform type is "known" with support from above (for example, reports) or with support from below (for example, data), then the facts about the platform type can serve as a model. From this model, we can hypothesize the platform's range of speed, its sound-producing machinery and the machinery's acoustic sig.. ·e, the platform's travel patterns, and so on. Pieces of data that can support the model-based hypothesis are sought in the signal data. As more supporting evidence is found, confidence in pursuing the model is increased. In this sense, the model serves as a constraint in the search process.

Additional Notes

1. Major differences in the design of the HEARSAY-II and HASP systems are summarized below:

> A knowledge source was written as procedures in HEARSAY and as a set of rules in HASP.

> Each knowledge source in HEARSAY consisted of two procedures: the condition part and the action part. In HASP the precondition part was a list of tokens, and the precondition and the action parts were in one module.

> In selecting a focus of attention, HEARSAY was concerned with selecting the next knowledge source to execute, and HASP was concerned with selecting the next solution island to pursue.

> HEARSAY used a central scheduler to select its focus of attention; HASP partitioned the scheduling task and used a hierarchy of control knowledge sources to select the focus of attention.

> In HEARSAY a subset of knowledge sources was chosen for execution from a list of all applicable (invoked) knowledge sources. HASP executed all knowledge sources applicable to a focused event. However, not all the changes to the blackboard became focused events.

> HASP was designed to interpret continuous, multiple streams of data. HEARSAY interpreted single speech utterance.

2. The hierarchical control in HASP was an attempt to separate the domain-specific knowledge from knowledge about the application of that knowledge. It was the first attempt at such an organization and was rather simplistic. In the CRYSALIS system (described next), the hierarchy of control knowledge sources was organized differently.

3. In HASP the control-related information was made globally accessible. It was also decided to represent control functions in rule form. The grouping of control related rules into control knowledge sources was an obvious next step. However, by not integrating the control information into the blackboard structure, the control rules had to be expressed and processed differently from the domain knowledge sources. The BB-1 system [18] corrects this awkward representation problem.

3.3. CRYSALIS
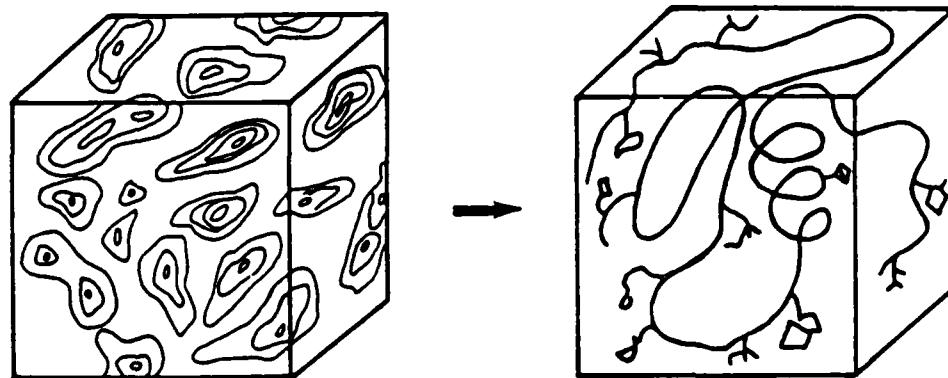
Figure 3-8:   The CRYSALIS Task

The CRYSALIS system contributed to the repertoire of blackboard system designs in two ways. First, it introduced the use of multiple hierarchies on the blackboard, the *blackboard panels*. Second, it addressed the control problem from the perspective of rule-based systems.[54]

The CRYSALIS project began in the spring of 1976 under the sponsorship of the National Science Foundation. It was a joint project between protein crystallographers at the University of California at San Diego and computer scientists int the Heuristic Programming Project at Stanford University. It was undertaken because the computer scientist thought that a blackboard approach "appeared to be appropriate" for this difficult task. The objective of the project was to build a system that determined the structures of proteins given their amino acid sequence and X-ray diffraction data for the protein crystals. The project did not reach its goal of building a system to construct complete stereo models of proteins. It did, however, succeed in a few cases in mapping more than 75% of the amino acid residues in the data. (As is seen later, this is equivalent to finding partial solutions on the middle level of the blackboard hierarchy.) Basically, this was a problem that usually took crystallographers months to solve and proved to be too difficult to solve completely. In retrospect, one can argue that this

---

[54]In this sense, the intellectual lineage of the control component of the CRYSALIS system is closely tied to MYCIN-like systems and applications written in OPS. One of the the major control issues in rule-based systems is to separate control information from domain rules and thus to make explicit the implied or built-in sequence of rule executions.

problem violated many of the criteria for choosing appropriate application problems: lack of experts from whom knowledge could be extracted and tested; almost no theoretical knowledge about the relation between the structures and the functions of proteins; and an impoverished state of knowledge about reasoning in three-space. Nonetheless, the CRYSALIS system did solve a significant part of the application problem and did contribute to the evolution of blackboard systems.

The Task

The task of the CRYSALIS system was to infer the three-dimensional structure of protein molecules. A protein structure was derived from an interpretation of the electron density map of the protein, which, in turn, is derived from X-ray diffraction data gathered from the crystallized protein. Traditionally, the protein crystallographer represents the explanation of the electron density map in a ball-and-stick molecular model fashioned from metal parts. These parts are strung together to form a model that conforms to the density map and is also consistent with protein chemistry and stereochemical constraints.

The *electron density map* is derived from diffraction patterns produced by placing a protein crystal in an X-ray beam. It records the density of electrons in the protein molecule sampled at various points on a three-dimensional lattice. The resolution of the electron density map is typically poor, and the locations of individual atoms are generally not identifiable.

In addition to the electron density map, CRYSALIS was provided with the amino acid sequence of the protein. These data could also be errorful, with the sequence being incomplete or out of order. Nevertheless, given the amino acid sequence, the problem is narrowed to a task of determining the folding of the chain of amino acid residues and peptide bonds consistent with the distribution of the electron density.

Other data were available to CRYSALIS that were not used directly by human model builders. They were produced by mathematical algorithms that abstracted (or reduced) the electron density data by keying in on different features of the density data (the equivalent of the low-level signal-processing algorithms used in the speech- and other signal-understanding tasks). One reduced data set consisted of density peaks above some threshold and their locations; another consisted of connected peaks and regions called *skeletons*; and a third consisted of segments of the molecular skeleton. In some sense, these data were pieces of solutions produced by three different knowledge sources and were useful as intermediate solutions on the blackboard. However, the initial attempt to construct a hierarchy that included the various data and a target molecular model proved unsatisfactory. First, the data representation did not integrate well with the abstraction hierarchy of the molecular model, which consisted of atom,

superatom, and secondary-structure levels. Second, the algorithms to generate the intermediate data were not designed to be used incrementally; that is, they could only work on the data for the entire molecule, not on small regions of the data. In order to organize all the data in a rational way, two hierarchical data structures were created for the blackboard: one to represent the bits and pieces of stereo structures conjectured during problem solving and the other to hold the data produced by the mathematical algorithms (see Figure 3-9).

The Blackboard Structure

The blackboard contained two abstraction hierarchies called *blackboard panels*. The density panel contained four levels: the raw electron density map, the peaks, the skeleton, and segments. The information on this panel was produced by signal-processing algorithms prior to the interpretation process. The hypothesis panel contained atom, superatom (amino acid residues and peptides), and stereotype (for example, alpha-helix) levels. The objective was to place each atom of the protein molecule in the three-space represented on the hypothesis panel. The solutions were built by generating partial solutions on the hypothesis panel derived from data at any level in the density panel.
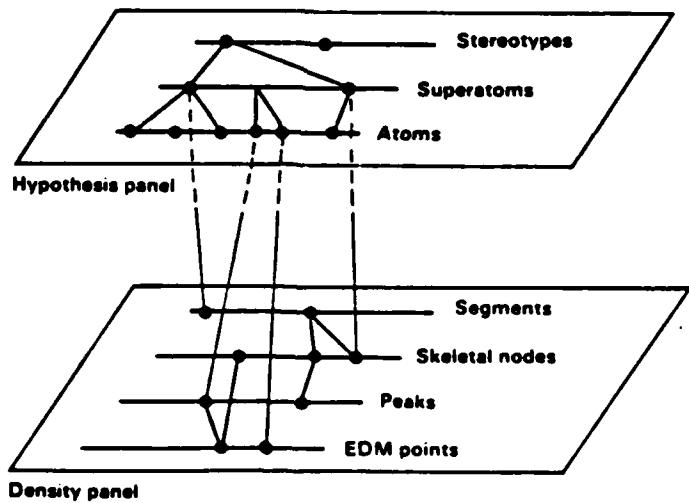

Figure 3-9:   The CRYSALIS Blackboard Panels

The Knowledge Source Structure

Each knowledge source consisted of a set of rules. Unlike the knowledge sources we have seen thus far, there were no preconditions associated with the knowledge sources. A precondition of a knowledge source served to inform the control module when it had something to contribute during the problem-solving process. The knowledge sources in CRYSALIS were not designed

to be self-selecting.

Control

The CRYSALIS system used a three-tiered control structure. All the control modules were uniformly represented as knowledge sources. The overall control of the system was assigned to the strategy-level knowledge source. A set of strategy rules governed the choice of the next task to be performed on particular regions of the blackboard. A task was represented as a task level knowledge source. Rules in the task knowledge source decided which object level knowledge sources to execute within the context of a given strategy.

The strategy knowledge source had access to a summary of the solution state called a feature list, which recorded the state of the solution by regions. The strategy knowledge source decided upon which region to work and, based on the characteristic features of that region, selected and executed a task level knowledge source. The selected task knowledge source executed a sequence of object-level knowledge sources based on the recent changes in the chosen region as recorded on the event list. The event list contained a list of changes made on the hypothesis panel (see Figure 3-10). After the execution of the object-level knowledge sources, control returned to the task knowledge source. The task knowledge source updated the feature list at this point and executed another sequence of object-level knowledge sources if the situation warranted. After the task-level knowledge source was finished, it returned control to the strategy knowledge source. The strategy knowledge source selected the next region on which to work and the appropriate task-level knowledge sources to reinitiate the processing. Only the object-level knowledge sources were allowed to modify the hypothesis panel, and no modification were made to the density panel.

One can view the organization of the control component in one of two ways: (1) as a nested activation of the knowledge sources, or (2) as an organization in which the precondition of each knowledge source was held within its immediate higher-level knowledge source. The higher level knowledge source acted as the manager of the lower-level knowledge sources. In either case, the opportunistic application of the knowledge sources was less evident in CRYSALIS than in other systems. First, a large region was selected (by strategy), then a series of specific nodes within that region were selected (by task). Finally a series of predetermined knowledge sources were executed to process each selected node. A task knowledge source remained in control until all the possible processing in a given region was exhausted. Once processed, a region was never revisited. The focus of attention consisted of subdividing a given region to find a solution island to process next.

In summary, the basic actions of the control component were the following:

1. The strategy knowledge source focused on a region of the blackboard based on information in the feature list and executed the appropriate task knowledge source.

2. The task knowledge source selected a specific place in the region and used it as a context for a sequence of object-level knowledge sources. The task knowledge source updated the feature list before returning control to the strategy knowledge source.

3. An object knowledge source modified the blackboard and returned control back to the task knowledge source.
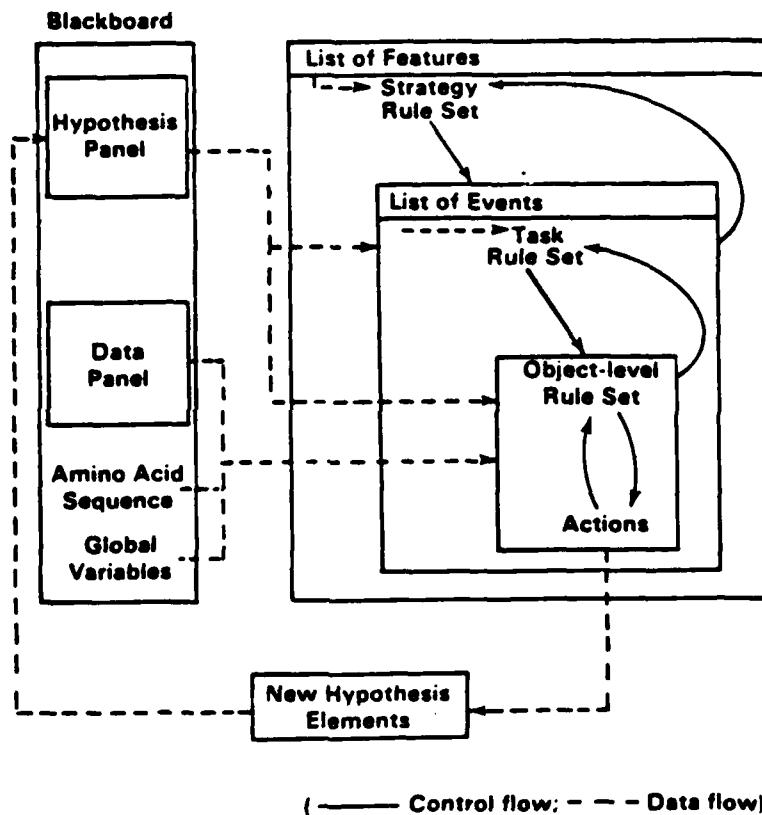

Figure 3-10:   CRYSALIS System Organization -- from [50]

Knowledge-Application Strategy

The behavior of the CRYSALIS system was strongly island driven, or more specifically, was directed at region growing on the blackboard. The following is a possible problem-solving scenario: Look for an electron-dense region that might indicate the presence of a tryptophan

element (a large, ring-containing amino acid). Look in the amino acid sequence for the occurrences of tryptophan. Look to see if an adjacent region might be one of the neighbors of a tryptophan element in the sequence. If the amino acid adjacent to the tryptophan matches the region in the close proximity of the tryptophan data, continue growing the region using the sequence as a guideline. When unable to continue with the region-growing process, look for another region to grow.

Given this problem-solving scenario, one can see the appropriateness of the CRYSALIS control scheme. The task-level knowledge sources, with names such as point-to-point-trace, outward-trace, split-group-toehold, and so on, knew which object-level knowledge sources to call in order to accomplish their goals (in the scenario, the goal of the task knowlege source is to extend the hypothesized region). The strategy knowledge source moved from one region to another, with each region demanding possibly different task goals. This type of nested processing was reflected in the hierarchy of the knowledge sources.

Within a region of interest, the selection of which node to process next was opportunistic. This was the only place that opportunism was exercised. The region selection followed the shape of the skeleton. The selection of the knowledge sources, both the Task knowledge sources and object knowledge sources within each task knowledge source, was built in.

Additional Notes

1. The quality of the density map affected the reliability of the knowledge sources. Each knowledge source had associated with it weights that could be adjusted to reflect the data quality. Furthermore, the rules were weighted according to an importance criterion. The weight on a hypothesis was a combination of the weights that reflected the data quality and the importance of the rule which generated the hypothesis.

2. In HEARSAY-II and HASP, knowledge sources were self-selecting. That is, the precondition of knowledge sources determined whether the knowledge sources were appropriate in a given solution state. In CRYSALIS there was no counterpart to the precondition of knowledge sources.[55]

---

[55] One wonders if CRYSALIS is truly a blackboard system in a strict sense because it violates one part of the definitions of the blackboard model, *knowledge sources respond to changes on the blackboard.* As mentioned earlier, the control component of blackboard systems has disparate designs. However, the knowledge sources should, in order to maintain their independence, indicate the condition under which they can contribute to the problem solving process. Because the knowledge sources in the CRYSALIS system were not designed so, I feel that the CRYSALIS control is a hybrid between a blackboard system and a rule-based system. It is a blackboard-like system.

3. As in HASP, the hypothesis panel in CRYSALIS was called the Current Best Hypothesis, and the nodes in the panel were called the hypothesis elements. In HASP the CBH represented the situation board created, updated, and used by the analysts for further interpretation of the signal data. In CRYSALIS the CBH represented the partial protein model built up to any given point in the model-building process. The hypothesis network on the blackboard represents a network of partial solutions, which is not necessarily the same as intermediate results. Whereas intermediate results often cannot stand on their own in the middle of a problem-solving process, partial solutions are often meaningful and useful on their own. Thus, if the CRYSALIS processing were to be interrupted and the hypothesis panel examined, there would be solution islands that are acceptable solutions.

3.4. TRICERO
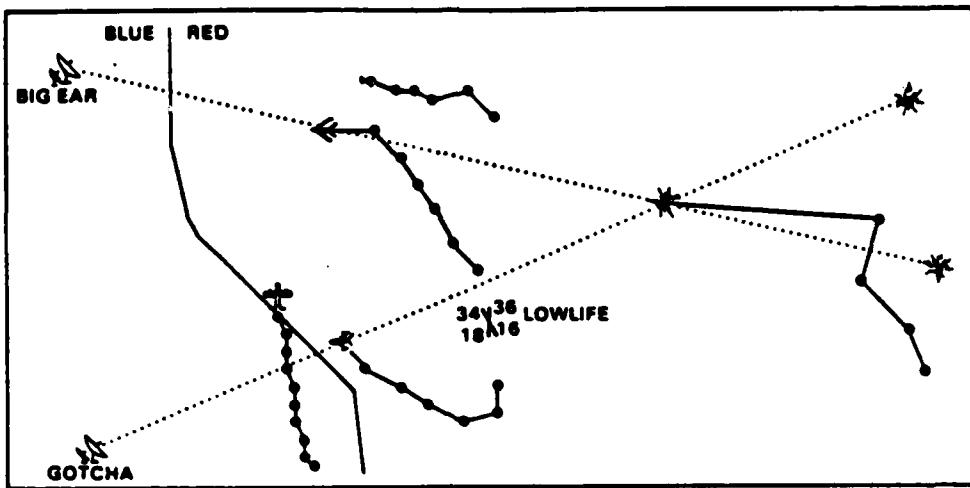
Figure 3-11:   The TRICERO Task

The TRICERO system represents an extension of the blackboard system into the area of distributed computing.[56] There are many possible ways to design a blackboard system that utilize multiple, communicating computers. To design a multi-processor blackboard system, either the blackboard model or the blackboard framework can be used as a design foundation. What is chosen as the starting point will have a significant effect on the nature of the concurrency in the resulting system.[57]

Sevaral possible ways exist for using multiple processors. First is to partition the solution space on the blackboard into loosely coupled regions (for example, sub-regions of the ocean, parts of a sentence, pieces of the protein structure, and so on). For each of these partitions, create a

------

[56]The TRICERO system was designed by Harold Brown of the Knowledge Systems Laboratory at Stanford University. It was built by programmers at ESL and Teknowledge. The TRICERO system was written using the AGE skeletal system [35]. The distributed-system aspects of TRICERO were simulated.

[57]Two parallel blackboard systems are currently being built at the Heuristic Programming Project. The design of one system is based on a fresh interpretation of the blackboard model. The system is designed to work within the context of a large number (100s to 1000s) of processor-memory pairs with high-bandwidth communication. [40] In the other system, designed as an extension to a serial skeletal system, parallel constructs are made available to the user. The design of this system is targeted for multi-processor, shared-memory systems. [2]

copy of a blackboard system. For example, in HASP one might have a complete blackboard system for each sensor array. Because the arrays have overlapping coverage, the systems would have to coordinate their problem-solving activities. A system will notify an "adjacent" system if a platform is moving into that system's area, for example. In other words, the application problem can be partitioned into loosely-coupled subproblems that need coordination. Research on this type of systems is being conducted at the University of Massachusetts under the direction of Victor Lesser [26]. A second way to use the blackboard data in a shared memory and distribute the knowledge sources on different processors (see [2] and [7] for examples). This distribution results in the parallel execution of the knowledge sources. If the knowledge sources are represented as rules, their condition parts can be evaluated in parallel. The action parts can also be executed concurrently with the evaluation of the condition parts in a pipeline fashion. The PSM project at Carnegie-Mellon University is targeted as a parallel rule execution system [14]. Third, a more direct use of multiple processors can be accomplished by partitioning the problem into independent subproblems, where each subproblem is solved on a separate processor. For example, in the interpretation and fusion of multiple types of data, each type of data might be interpreted on different systems. Each system will have a different set of knowledge sources and a different blackboard organization. The results from the data analysis systems will be fused by another blackboard system. The TRICERO system is an example of this type of system.[58]

The Task

The objective of the TRICERO system is to monitor a region of airspace for aircraft activities. The system consists of three subsystems organized in an hierarchy (two levels at this point), much like the human management organization for which the system was built (see Figure 3-12). On the lower level are the ELINT and COMINT subsystems that respectively interpret passive radar and voice communication data. The correlation subsystem that integrates the reports from ELINT and COMINT and other data resides at a higher level. This hierarchical organization of blackboard system emulates the various activities involved in signal understanding. These activities are signal detection, parameter estimation, collection analysis, correlation, and overall interpretation. As one progresses from one activity to another, information in the data is abstracted and reduced. TRICERO analyzed two types of collection

---

[58]The problems of designing and building blackboard systems capable of concurrent problem solving, distributed problem solving, and parallel computations are distinct from those of serial blackboard systems and are not discussed in this document. The TRICERO system is discussed here, because it does not fall into any of the earlier categories. It is a variant of a distributed computing system that can be considered a direct extension of the serial systems. See [26] for distinctions between distributed-processing systems and distributed problem-solving systems.

data and correlated the analyzed data. Each data type was analyzed independently using different blackboard data organizations and different knowledge sources (see Figure 3-12).


(— data/messages)
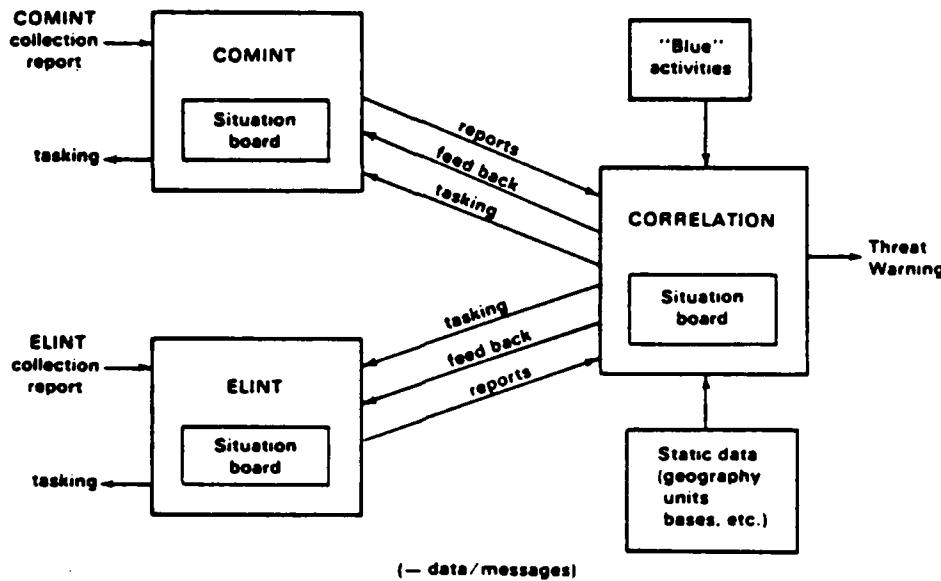
Figure 3-12:   A Distributed Blackboard System, the TRICERO Control

The Blackboard Structure

The ELINT blackboard consisted of three levels: observation, emitter, and cluster. The input data arrived at the observation level. These data were tagged with the collection time and the site at which they were collected. Each node on the emitter level kept a history of detections from a site having the same identification tag. The history represented radar emissions believed to be emanating from one source. The identification tag could be in error, whereby different sources could have the same identification tag, or one source could have multiple tags. The radar emissions detected at different sites were merged into an hypothetical platform (ui a number of platforms "seen" as one platform) on the cluster level. Each level used descriptive vocabulary appropriate to that level: the platform types and speed history on the cluster level and the collection site and signal quality on the observation level, for example. The blackboard data structure in the COMINT and correlation subsystems were structured in similar fashions using abstraction levels appropriate to interpreting their data.

The Knowledge Source Structure

The knowledge sources were structured according to the specification in the AGE [35] skeletal system. Each knowledge source had a precondition part and an action part. The precondition part was a list of tokens, each representing a type of change that could be made on the blackboard. The action part consisted of a set of rules. The rules in each knowledge source could be processed as a multiple hit, in which all rules whose condition sides were satisfied were executed, or as a single hit, in which only the first rule whose condition side was satisfied was executed. There was no conflict-resolution process of the type found in OPS-based systems.

Control

Each of the independent subsystems in the TRICERO system used a subset of control components available in AGE. A globally accessible event list recorded the changes to the blackboard. At each control cycle, one event on an event list was chosen as a focus of attention. The choice of the event on which to focus was based on a predetermined priority of event types. Once an event (an event type and a node) was selected, it was matched against the event-type tokens in the precondition of the knowledge sources. Those knowledge sources whose preconditions contained the event-type token matching the focused event type were executed according to a predetermined priority of knowledge sources.

The TRICERO system augmented the AGE control component to handle the communication among the three subsystems. Each subsystem could send messages to designated subsystems. The receipt of a message by a subsystem was treated as an event focused on a special node on the blackboard. This construct allowed the subsystems to treat reports from other subsystems just like any other event.

The basic actions of the control component can be described in two parts:

Between subsystems

1. The simulation of the distributed computation consisted of round-robin execution of the three subsystems -- ELINT, COMINT, and correlation.

2. Each subsystem sent report messages to designated subsystems. The receipt of a message was treated as an event with appropriate modification of the recipient's blackboard and event list.

Within a subsystem

1. A control module selected a focus event using a list of event priorities. An event

61

contained information about the event type of the change made to the blackboard, that is, the node on which the change was made, the knowledge source and the specific rule that made the change, and the actual change.

2. Based on the focused event, knowledge sources whose precondition list contained the event were chosen for execution.

.3: The rules in the activated knowledge sources were evaluated and executed according to the rule-processing method associated with the knowledge source. Modifications to the blackboard by the rules were events and caused the event to be put on the event list.

Knowledge-Application Strategy

Most of the knowledge sources engaged in bottom-up processing. They combined information on one level to generate a hypothesis on a higher level. The reports from the correlation subsystem to ELINT and COMINT dealt primarily with information on the higher level (for example, platform identification) that overrode the analysis done by the lower level subsystem. In such cases, the processing in these subsystems became top down. The reports from ELINT and COMINT were treated as input to the higher-level correlation subsystem.

Additional Notes

1. The partitioning of the overall task into subsystems in TRICERO was accomplished by assigning the analysis of the abstract information to the correlation subsystem and the analysis of information closer to signal data to ELINT and COMINT. As mentioned in Part One, the knowledge sources that span the various levels of the blackboard hierarchy are logically independent. Thus, the need for coordination among the subsystems is substantially reduced when the problem is partitioned into subsystems along carefully chosen levels of analysis.

2. As with the other systems described, the TRICERO data were noisy and the knowledge sources uncertain. The radar data, for example, contained "ghosts," detections of nonexisting objects. The ELINT subsystem handled the existence of this type of error by delaying the analysis until several contiguous detections had occurred. By doing so, it avoided the creation of hypothesis nodes that later needed to be deleted.

The issues relating to the deletions of nodes on the blackboard are quite complex. Suppose in TRICERO that a node on the cluster level (an object that represents a platform or a group of platforms) is to be deleted. What does it mean? Has the platform disappeared? Unless it somehow disintegrated, a platform cannot disappear into thin air. Was there an error in interpreting the radar data to begin with? Often, there are "ghost tracks," a characteristic of

which is that the tracks disappear after a short duration. However, suppose the platform disappearance was not due to ghost tracks but to an error in reasoning. Unraveling the reasoning steps that led to the hypothesis (backtracking) and retrying often do not help. The system does not know any more than it did when the erroneous hypothesis was generated. Suppose the platform node is just deleted. What do we do about the network of evidence that supports the existence of the platform? Unfortunately, there is no systematic way of handling node deletions. In HASP the nodes were never deleted. The nodes in error were, ignored and analysis continued ignoring past errors. In TRICERO node creation was delayed until there was strong supporting evidence for the existence of an object represented by the node. When an error occurred, the hypothesis network was restructured according to domain heuristics.

3. In TRICERO the confidence assigned to the hypothesis elements was expressed in a symbolic form. The vocabulary expressing the confidence consisted of "possible," "probable," "positive," and "was positive." The confidence level was changed according to heuristic criteria.

4. TRICERO was one of the first blackboard systems implemented on a computer system with a bit-map display (see Figure 3-11 for a display output). The situation board, symbolically represented on the blackboard of the correlation subsystem, was displayed in terms of objects in an airspace and the objects' past behavior. The graphic-display routines were written as procedural knowledge sources and were executed when certain events (changes on the blackboard) occurred that warranted display updates. There might be some argument about the conceptual consistency of this approach because interfacing is usually not considered a part of problem solving. However, this engineering solution that integrated the display routines with the problem solving components worked very well. An effective display interface requires knowledge about what is appropriate to display when. A knowledge-based control of displays and display updates is easily implemented using the knowledge source organization.

3.5. Other Blackboard Systems
The four application systems discussed thus far transformed signal data into symbolic forms "natural" to the task domain. The signal-to-symbol transformation occurred for the purposes of understanding the context in which the signals were present. There are other blackboard systems that deal with similar application problems. Unfortunately, many of these systems are either proprietary or classified. The descriptions of these systems lack technical details, and we were not able to include them for discussion. We have included references to articles describing some of these systems -- see [23], [29], and [48].

We now turn our attention to two blackboard application systems that have been built to address different types of tasks. A brief description of the task is followed by some notable features of these systems.

3.5.1. OPM

The OPM system differs from the systems described so far in that it is a simulation of a model of human cognitive processes in planning. The cognitive model was reflected in the architecture of the OPM system. Instead of generating a new plan, it replicated the planning process of human subjects. The fact that the system seemed to successfully model many subjects' planning processes indicated the validity of the model. It also attested "to the utility of the blackboard model as a general model of cognition." [20]

In addition to the cognitive-modeling aspect, OPM demonstrated the generality of the blackboard model in the kinds of tasks it could address. The applications we have discussed so far dealt with interpretation tasks that are basically analytic in nature. The processing is primarily bottom up. The planning task is primarily generative in nature. It starts at the top with a goal to be achieved, and the planning process then produces lower-level sequences of actions to be performed. As we saw in both HEARSAY-II and HASP, top-down strategies were combined with bottom-up strategies to interpret noisy data, but nonetheless, the interpretation process was strongly data driven. The utility of the blackboard model for planning tasks opened up the possibility of building blackboard systems for many different classes of application problems.

The Task

The objective of OPM was to simulate human errand-planning protocols. "The planner begins with a list of desired errands and a map of a town in which she or he must perform the errands. The errands differ implicitly in importance and the amount of time required to perform them. The planner also has prescribed starting and finishing times and locations. Ordinarily, the available time does not permit performance of all of the errands. Given these requirements, the planner decides which errands to perform, how much time to allocate for each errand, in what order to perform the errands, and by what routes to travel between successive errands." [20]

The Blackboard Structure

The blackboard was partitioned into five planning panels, called planes, containing conceptually different categories of decisions. Each panel contained several levels of abstraction found in the planning space. The five panels were (see Figure 3-13):

1. Meta-Plan: Decisions on this panel indicated what the planner intended to do during the planning process. For example, on the policies level, a knowledge source specified general criteria to impose on the problem solution, such as "the plan must be efficient" or "minimize certain risks."

2. **Plan Abstraction**: Decisions on this panel characterized desired attributes of potential plans. These abstract decisions served as heuristic aids to the planning process, suggesting potentially useful qualities of the planned actions.

3. **Knowledge Base**: This panel recorded observations and computations about relationships in the world that the planner generated while planning. This knowledge supported two types of planning functions--the analysis of the current state of affairs and the analysis of the likely consequences of hypothesized actions. For example, at the errand level, the planner might have computed the time required to perform all of the currently intended errands in order that the planner might evaluate the plan's gross feasibility.

4. **Plan**: Decisions on this panel indicated actions that the planner actually intended to take. Decisions at each level within this panel specified a more refined plan than those at the adjacent higher level. For example, the outcomes level indicated what the planner intended to accomplish by executing the final plan, whereas the procedures level specified specific sequences of actions (errands).

5. **Executive**: This panel contained information related to control. The knowledge sources on this panel decided which of the invoked specialists were to be executed. The decisions were based on information on the different levels representing different types of "executive decisions." For example, priority decisions indicated a preference for allocating processing activity to certain areas of the planning blackboard before others. Schedule decisions indicated which of the invoked specialists satisfying higher-level decisions to execute next.

The Knowledge Source Structure

The knowledge sources were called *specialists*. Each specialist consisted of two components, condition and action, as in HEARSAY-II. The condition part consisted of a trigger and a test. The *trigger* provided a quick preliminary test of a specialist's relevance for any focused node. The *test* specified all other prerequisites of applicability. For a given focus node, the triggers of all specialists were checked. Test parts were evaluated for those specialists whose triggers were satisfied. A specialist became "invoked," as in HEARSAY-II, when both the trigger and the test parts of the condition were satisfied.

The specialists in OPM were written as procedures, as in HEARSAY-II. However, the procedures were much smaller and represented a smaller grain of knowledge.

Control

The OPM system had four global data structures: a map on which the errands occurred, the blackboard that contained the partial solutions, an agenda that held a list of invoked specialists
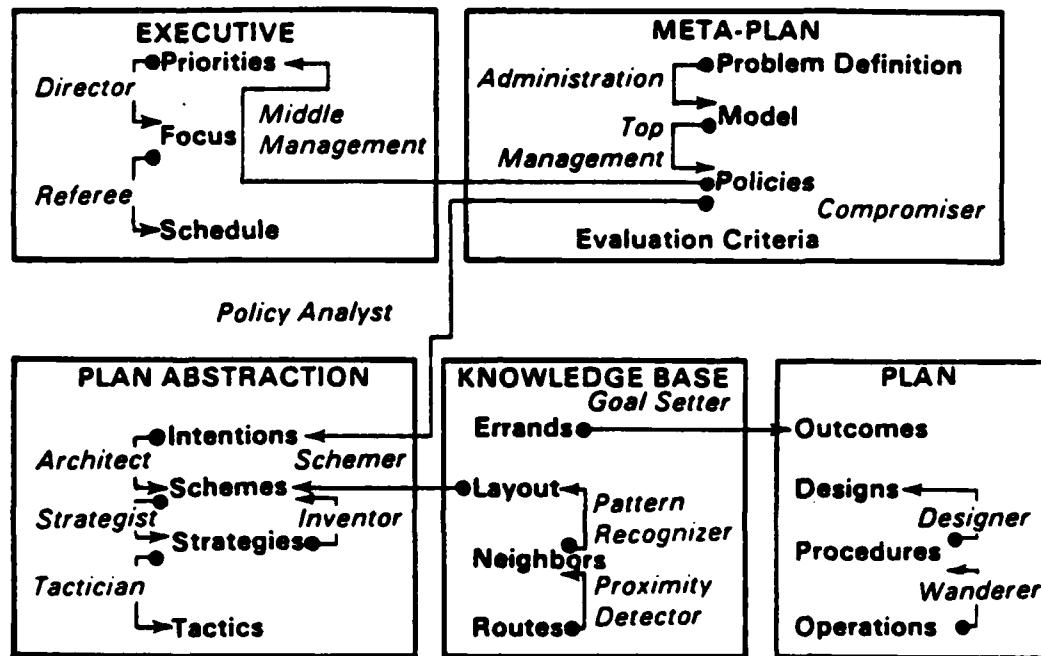
Figure 3-13: OPM Blackboard and Knowledge Sources

which needed to be scheduled for execution, and an event list that recorded the history of changes to the blackboard.

The basic actions of the control component consisted of three phases that were repeated:

1. During the invocation phase, the test part of the specialists on the agenda was evaluated. Specialists whose tests were satisfied became "invoked". The program terminated when there were no invoked specialists.

2. In the scheduling phase, one of the invoked specialists was recommended for execution. The basis of the recommendation was recency of invocation and current focus. The focus node was the most recently added or modified node on the blackboard. A specialist was chosen whose action, if executed, would occur in the region of the focused node.

3. In the execution phase the scheduled specialist was executed. The program immediately evaluated the trigger of all specialists against the *focus node* (the one just changed) and added those specialists whose triggers were satisfied on the agenda.

Knowledge-Application Strategy

The objective of the system was to enable the simulation of diverse problem solving behavior exhibited by human subjects while planning. Some example behaviors are describe in [20].

As can been seen from the design of the control modules, however, the basic strategy was, for psychological reasons, to follow up the most recent actions and in the geographic proximity of these actions. The psychological reason is not explained. However, one can surmise that at least for the errands task if a person decides to do an errand in one place, then that person will do all the errands in the same area.

Additional Notes

1. The OPM design reflects the first step taken to separate and make independent the problem of control. A scheduler of the HEARSAY-II variety was encoded as knowledge sources. Each of these control knowledge sources had a specialized scheduling policy, for example, go to the next closest place or do the next most important errand. The information needed and generated by the control knowledge sources was stored on a special blackboard panel, the executive panel. The executive panel was isolated from the other panels in that changes on the executive panel were not treated as events which affected the triggering of other knowledge sources; changes only affected the selection of the knowledge sources to be executed.

2. The map data resided outside of the blackboard data structure. It might have been interesting to organize the map on a blackboard panel in a similar manner to the CRYSALIS density panel. One could then have modeled the generation of abstract plans and strategies from abstracted maps.

3.5.2. Scene Understanding
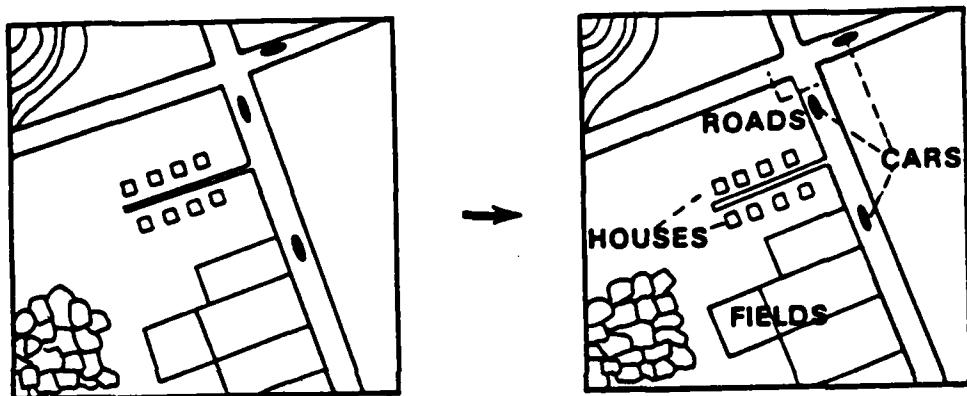
Figure 3-14:   The Scene Understanding Task

As mentioned in 2.2, one stimulus for the HEARSAY project was a desire to integrate syntactic and semantic information into the understanding of speech utterances. The notion of semantically driven vision programs had also existed for some time. A program to interpret complex aerial photographs, developed by Makoto Nagao and Takashi Mastuyama, was the first to address that problem using the blackboard approach.

The Task

Given a large aerial photograph of a complex suburban area taken at low altitude, the program is to identify and label objects in the photograph (see Figure 3-14). The variety of objects to be identified includes, among other things, cars, houses, rivers, and roads.

The Blackboard Structure

The blackboard was organized into an abstraction hierarchy consisting of elementary regions, characteristic regions, and object levels (see Figure 3-15). The lowest level, the elementary regions level, contained regions segmented according to multispectral properties. The attributes of each elementary region were its average grey level, size, location, and basic features, together with pointers to the digitized picture indicating its relative position. A combination of elementary regions formed an object on the characteristic regions level. On this level, seven characteristic features were extracted: large homogeneous regions, elongated regions, shadow regions, shadow-making regions, water, vegetation, and high contrast-texture regions. The objects from this level were classified into one of the domain objects the system knew about

-- cars, houses, rivers, roads, crop fields, and so on.
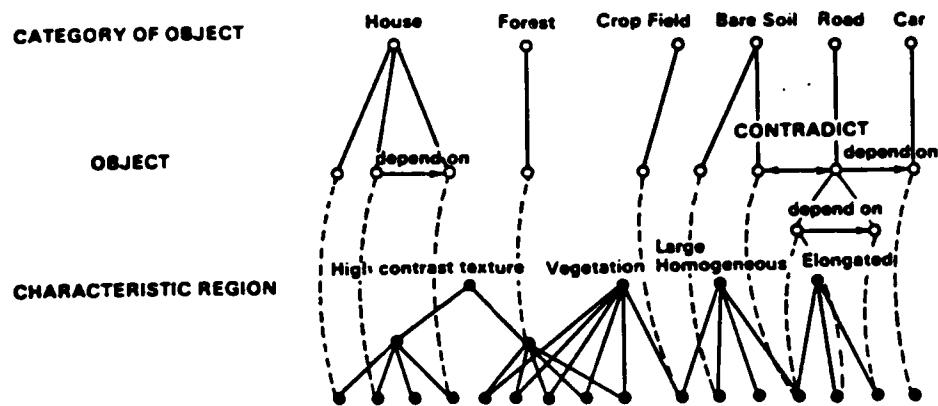
Figure 3-15:   Blackboard and Knowledge Sources for Scene Understanding

The Knowledge Source Structure

A knowledge source was represented as a single rule. Thus, the condition part of the rule served as the precondition of knowledge sources. Each rule held enough knowledge to recognize one object. Because there are many different ways to recognize an object, there were multiple rules (knowledge sources) for the recognition of a specific type of object.

Control

All knowledge sources (rules) whose condition sides were satisfied by the data in the blackboard were executed, as in HASP. Thus, no complex scheduling module was needed. However, because knowledge sources might interpret a region differently (for example, as a crop field or grassland), the system incorporated a mechanism to resolve this type of contradiction. Called the *conflict-resolution* mechanism, it calculated the "reliability value" of each region interpretation and retained only the most reliable interpretations.

Because applying each object-detection knowledge source to every region would be computationally expensive, the knowledge sources were applied only to those regions with certain characteristics. These characteristics were precisely the attributes of the objects on the characteristic-region level. Thus, the preconditions of the knowledge sources served as a filter that kept the knowledge sources from processing each region. The authors called this process the focusing mechanism. It emulates the way in which a human first globally surveys a scene

to find prominent features that attract interest before doing detailed analysis.

Knowledge-Application Strategy

Each object category in the system was specified *a priori* and had associated with it knowledge sources that could recognize or reject an object in the category. For example, there was a "house expert" that, given a region, could recognize it as a house or could reject the possibility of it being a house. The knowledge sources were further specialized into data-driven knowledge sources and model-driven knowledge sources. Data-driven knowledge sources were capable of combining objects on the characteristic-region level into identifiable domain objects. The model-driven knowledge sources interpreted regions on the characteristic-region level based on already identified objects. For example, cars could be easily identified once a road had been identified.

The object-recognition process thus worked both bottom up and top down. Aside from the model-driven aspect, the system employed another form of top-down processing not found in the systems discussed thus far. If the object-detection knowledge sources were unable to detect an object (that is, unable to label a region as an object), the system reapplied segmentation subsystems in order to split the region or to merge it with adjoining regions. The newly formed region was then processed again by object-detection knowledge sources.

Additional Notes

1. The integration of symbolic reasoning and numeric (or algorithmic) computation is known to be important, especially the feedback of information from the symbolic side to the numeric side. A push from the symbolic side basically amounts to top-down processing, whether it be for model-driven analysis or for a goal-directed analysis. Nagao and Matsuyama's system is the first documented system we could find that actually accomplished the integration. In blackboard systems, the integration of symbolic and numeric processes would appear simple -- one only needs to treat a numeric algorithm as another knowledge source. One reason for more systems not having this integration, for example in HASP or CRYSALIS, is that the integration must be planned from the beginning and the mathematical algorithms written to fit the plan. Neither usually happens.[59] For example, the algorithms that would have been useful in CRYSALIS were not written to serve as knowledge sources. One requirement for a

---

[59]This is not surprising. There is a "cultural" gap between people who build numeric algorithms and those involved in symbolic reasoning. Each group would like to solve a given problem within their own discipline. However, there is a great deal to be gained by combining the strength of both sides.

knowledge source is its ability to deal with partial solutions or data. Most algorithms are designed to process whole data. In addition, for the algorithms to be used effectively, other knowledge sources must be able to set parameters in the algorithm. For example, a knowledge source might ask an algorithm to increase the data-sampling rate to see if the information content in the algorithms' output can be increased or improved. The effective utilization of numeric algorithms within a blackboard system is a knowledge-based task.

3.6. Summary

A definition of the blackboard model and a summary of blackboard system design, in the form of the blackboard framework were provided in Section 1. Section 3 reviewed some of the older application systems. Although a problem-solving model can help in the general organization of domain knowledge and reasoning strategy, the blueprint of the architecture is drawn by the characteristics of the specific task at hand. Details of the task determine the specific choice of knowledge representation and reasoning methods. It is possible, therefore, that there are as many blackboard architectures as there are applications. Figure 3-16, which summarizes a small set of blackboard systems, indicates variations in the characteristics of signal-understanding problems and variations in the blackboard systems designed to solve the problems. To what degree the differences in the design, especially in the design of the control component, can be attributed to the differences in the problem is hard to determine at this point. At the same time, the figure indicates that complex problems can be solved using the blackboard problem-solving organization and that the basic organization allows for a wide range of variations in system designs.

Application Characteristics

	HEARSAY-II	HASP	CRYSALIS	TRICERO	OPM	ACAP*
Task	speech understanding	sonar signal understanding and information fusion	electron-density map interpretation	military signal understanding and information fusion	errand planning	image understanding
Continuous data? (evolutionary solution)	no	yes	no	yes	no	no
Signal-to-noise ratio	moderate/high	low/moderate	low	moderate/high	N.A.	moderate/high
Uses solution-space generator?	yes	no	no	no	no	no
Data reliability	moderate	moderate	high	moderate	high	moderate
Partial solution evaluation?	yes	no	yes	no	yes	no
Belief revision?	no	yes	no	yes	yes	yes
Multiple hypotheses?	yes	no	no	no	no	no

All the application tasks have (1) big factorable solution space, (2) interactive subproblems, (3) multiple lines of reasoning, (4) heterogeneous domain vocabulary, (5) no fixed sequence of subproblems to be solved. (Expert-system characteristics from Hayes-Roth, et al., 1983, p. 91)

System Characteristics

	HEARSAY-II	HASP	CRYSALIS	TRICERO	OPM	ACAP*
Number of panels and levels	1/7	1/6	2/3	3/4†	5/4†	1/5
Knowledge source condition (form)	procedure	event name	none (imbedded in control modules)	list of event names	procedure	condition part of a rule
Knowledge source body (form)	procedure	set of rules each KS	set of rules control KS	set of rules each KS	procedure	action part of a rule
What determines an event?	monitor				monitor	monitor
Focus of attention	a KS	a blackboard node	a blackboard node and a KS	a blackboard node	a KS	a KS and blackboard data
What determines the locus of attention?	scheduler	event manager	control KS	priority manager	scheduler	[predetermined]
Primary processing strategy	generate and test with constraints; bottom-up generation and test; top-down and middle-out constraint generation	inductive synthesis combined with model-derived expectations; bottom-up synthesis and top-down expectation generation	model-guided synthesis and region growing; model generation from auxiliary data and middle-out region growing	inductive synthesis and correlation; bottom-up synthesis and high level correlation	generate and test with refinement; top-down plan generation and bottom-up refinement	feature extraction and model-driven feature match; bottom-up feature extraction and top-down match

*Analysis of complex aerial photograph
†Average number of levels

Figure 16. Summary Characteristics of the Application Systems
[This table was developed by Harold Brown, Bob Engelmore, and Penny Nii]

4. Blackboard Systems from a Knowledge Engineering Perspective

There are many blackboard or blackboard-like systems being developed that still need to be analyzed.[60] Instead of a conclusion, this section contains one knowledge engineer's observations about the advantages and drawbacks of using a blackboard approach for building expert systems.

When should the use of a blackboard model be considered by a knowledge engineer? A general guideline is that a blackboard approach is useful for complex, ill-structured problems.

Complex Problems

Simon [46] defines a complex system as "one made up of a large numbers of parts that interact in a nonsimple way. In such systems, the whole is more than the sum of the parts, [in the sense that] given the properties of parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole." In order to understand complexity, we describe complex systems in terms of subsystems and relationships between the subsystems that are less complex. Often, this description takes the form of a hierarchy.

In software engineering, there are techniques and methodologies that foster hierarchic problem decomposition, and most complex programs are organized according to some form of hierarchy. *Usually, the hierarchy is organized along functional decompositions of the task to be performed.* This organization has the advantage of allowing common functions to be shared by many subsystems. In blackboard systems, a problem is decomposed to maximize the independence of the subsystems. Functionally, subsystems that generate and fill the solution space are separated from subsystems which determine their utility. All of these subsystems (knowledge sources) can be organized into a control hierarchy, but the emphasis is on limited interactions between the subsystems. This allows for maximal flexibility in the software development phases as well as during the problem-solving phase.

The medium of interactions among the subsystems (the blackboard) is also organized hierarchically. The hierarchical organization of the solution space on the blackboard has pragmatic advantages in designing a solution to a problem. First, the hierarchical structure allows for the integration of diverse concepts and associated vocabularies. For example, in HASP the concept of a "platform" was defined with properties such as type, speed, and location; the concept of "signal" was defined with properties such as frequency, intensity, and bandwidth. Second, the abstraction of information reduces the computational need in two

---

[60]We solicit builders of blackboard systems to provide information so the summary table can be expanded.

ways: (1) the manipulations of abstract entities involve manipulation of smaller numbers of entities than manipulations of their detailed counterparts and (2) abstractions can store information that would otherwise need to be recomputed from the detailed counterparts. Put another way, the hierarchical structure of the blackboard provides a favorable trade-off between storage space and computational time. For example, in CRYSALIS reasoning with an amino acid as an object in its own right is easier and faster than reasoning at the level of its atomic constituents, even though extra storage is needed to represent the amino-acid object.

Ill-Structured Problems

Ill-structured problems [31] are characterized by poorly defined goals and an absence of a predetermined decision path from the initial state to a goal. Often, there is a lack of well-defined criteria for determining whether a solution is acceptable.[61] Being ill structured is sometimes intrinsic to a problem; for example, sculpt a masterpiece. At other times, a problem is ill structured because it is ill defined or ill understood; for example, assess the merits of one's financial investments in light of the proposed tax reforms.[62] Although Newell's discussions about ill structured problems occur within the context of weak problem-solving methods, he notes that a human's ability to solve an ill-structured problem may be due to the problem solver's ability to "recognize the essential connection or form of the solution" or due to the fact that "the problem solver always [has] available some distinctions that apply to every situation." [31] In short, many ill-structured problems might be solved by applying knowledge, especially knowledge in the form of empirical associations, or expertise.

Many of the current expert systems deal quite well with ill-structured problems. What further aid can the blackboard approach provide? First, the blackboard approach requires no *a priori* determined reasoning path. Because ill-structured problems do not have a predetermined decision path to a solution, the selection of what to do next must be made while the problem is in the process of being solved. The incremental and opportunistic problem-solving approach in blackboard systems provides the capability to do precisely that. Second, from a knowledge engineering viewpoint, vague information and knowledge, which characterize ill-structured problems, need to be made concrete in the process of finding a solution to the problem. The blackboard model is an excellent tool for this knowledge engineering activity. (The blackboard

---

[61]In many expert systems, the acceptability of a solution is determined by a panel of human experts who might disagree among themselves.

[62]It appears that some problems are more ill structured than others. Art-making has traditionally been considered a very ill-structured task. But, AARON program makes "freehand" drawings. [5] It's interesting that AARON is a blackboard-like system.

model as a problem-formulation tool is discussed in the next section.) The blackboard approach is also an excellent tool for *exploratory programming,* a useful technique for developing complex and ill-structured problems and is discussed in the section on the use of the blackboard model as a development tool.

Although useful for many complex, ill-structured problems, blackboard systems are generally expensive to build and to run. It would be foolish to apply the blackboard approach when lower-cost methods will suffice. For example, classification problems [4] can in principle be solved using the blackboard method, but there are lower-cost approaches to the problem. Determining the appropriate problem-solving methodology for an application problem is itself a difficult problem and is not one of the topics of this paper. The reader is referred to [21], [17], and [52] for some guidance. Generally, the occurrence of some combination of the following characteristics in a problem makes it an appropriate candidate for the blackboard approach:

- A large solution space
- Noisy and unreliable data
- A variety of input data and a need to integrate diverse information
- The need for many independent or semi-independent pieces of knowledge to cooperate in forming a solution
- The need to use multiple reasoning methods (for example, backward and forward reasoning)
- The need for multiple lines of reasoning
- The need for an evolutionary solution

4.1. The Blackboard Model as a Problem-Formulation Tool

During the preliminary knowledge engineering phase, the goal of the knowledge engineer is to understand the task domain and the objectives of the proposed system. A knowledge engineer needs a set of conceptual models for organizing knowledge and reasoning. During the initial interactions with an expert, a knowledge engineer tries to find an appropriate conceptual model for the task while trying to understand the domain and the nature of the task. Often, the understanding of the task occurs with the help of a conceptual model. Because the information provided by the expert is rarely organized to fit a particular problem-solving model, the knowledge engineer initially needs a model with flexible methods for representing and applying pieces of knowledge.[63] Many of the issues faced by a knowledge engineer are the same as those faced by a traditional software engineer. As a knowledge-engineering tool, the blackboard

---

[63]Often, it is useful to use the blackboard model as a conceptual tool for understanding the task. Once the task is well understood, it can be reformulated for a simpler and less expensive problem-solving method.

approach is useful because it provides some organizational principles that are both powerful and flexible.

Partitioning the Knowledge

To make complex problems manageable, they often need to be decomposed into loosely coupled subproblems. As mentioned earlier, one useful decomposition tool is the partitioning of the solution space into a hierarchy. Hypothesizing the objects and their relationships is accomplished by applying knowledge. If the blackboard hierarchy is organized correctly, then the knowledge sources that operate between the levels in the hierarchy should be more or less self-contained. That is, a knowledge source should function much like a specialist, requiring little information from levels other than the one in which it is expert. For example, a lexicographer should not need much information from a phoneticist. Put another way, the system should have the behavioral characteristics of a nearly decomposable system as described by Simon [47] -- there should be less communication among the subsystems (knowledge source, in this case) than communication within each subsystem.

Separating knowledge and the uses of knowledge

A piece of knowledge can be used for many purposes. For example, knowledge about statistical methods can be used for processing speech, sonar, X-ray, radar, and visual signals. Whether a particular method, for example a least squares method, is useful depends on what it is to be used for, on the goals of the application problem and on the specific situation that arises while solving the problem. Thus, it is useful for a knowledge engineer to have a model that explicitly separates knowledge and the when, where, and how of applying the knowledge. The blackboard organization encourages the designer to make these separations. In addition, in a blackboard system, decisions as to when, where, and how a piece of knowledge is to be applied are made dynamically. The knowledge engineer can thus design a problem-solving strategy or set of strategies that best exploits the state of the solution. The separation also allows the knowledge to be expressed, at least in principle, in a "pure" form, unencumbered by information on how or when or where it is to be used. If the application of knowledge is in itself a complex task, then the blackboard model allows for the control to be encoded as knowledge sources that specialize in reasoning about knowledge application. This capability provides another dimension of organizational flexibility.

Errorful Data

Both noisy input data and unreliable knowledge can result in erroneous partial solutions on the blackboard. Traditionally, uncertainty in the knowledge has been solved by assigning credibility weights to the knowledge and to the information it generates. The blackboard approach provides an additional methodology for coping with errorful data. The basic idea is to establish independent evidential support and reasoning paths to a solution. Many different

kinds of data can be used to generate a hypothesis; for example, both intelligence reports and signal data were used in HASP. In addition, both HASP and TRICERO, whose input data were frequent observations of the same scene, exploited information redundancy in the data. In addition to the use of diverse sources of information, different expertise (competing knowledge sources) can be used. In HEARSAY-II both the syntactic-semantic hypothesizer and word-candidate generator were experts in generating the next word in a sentence. In HEARSAY-II, HASP, and other blackboard systems, hypothesized partial solutions had reasoned supports from above and supports from below to compensate for erroneous data and uncertain knowledge. The ability to use different resources, whether they be additional data or knowledge, to address the data error problem is important. The advantage of the blackboard approach is that organizationally, at least, it does not preclude the use of any of these resources.

4.2. The Blackboard Model as a System Development Tool

In traditional programming practices, system building consists of determining the requirements, designing a system, and implementing and testing a program. A programmer works from detailed system specifications. Almost all current software engineering techniques (for example, structured programming) are designed to ensure that the implementation strictly follows the specifications. The test phase consists of checking that the program performs according to the specifications. One aspect of dealing with complex and ill-structured problems is that there is often no *a priori* specification for a system. An additional difficulty with designing and implementing expert systems is that at least two parties are constantly shifting their points of view: the domain expert and the knowledge engineer. As the knowledge in the program accumulates and the problem becomes clearer, the knowledge engineer might find better ways to represent and process the knowledge. The resulting behavior of the program might inspire the expert to shift his view of the problem, creating further problems to solve for the knowledge engineer. Consequently, a knowledge engineer needs to engage in exploratory programming. Exploratory programming, as defined by Beau Sheil is "a conscious intertwining of system design and implementation."[64] [44] Waterman [51] notes that expert system building is accomplished in developmental stages ranging from research prototype to demonstration prototype to field prototype and so on until a fielded system is evolved. That is, expert systems are also developed incrementally. At each development stage, each of the system building phases (design, implement, and test) is repeated. The test phase is different from the traditional approach in that what is being tested is the overall acceptability of the behavior of the system and not the adherence to specification. At each development phase, the expert system might have to be redesigned and rebuilt.

---

[64]Sheil's article contains an excellent discussion on the necessity and the dimensions of exploratory programming.

In light of a need for exploratory and incremental system development, what is desired is a robust system organization that allows for modifications and additions at each development stage with minimum perturbation to extant structures and code. The HEARSAY-II and HASP systems went through various changes without major overhauls. In HEARSAY-II various configurations of knowledge sources and blackboard levels were tried, parallel constructs were added to run on multiprocessors, and various scheduling schemes were experimented with. The HASP system was converted from an interpreted system to a compiled system, the sonar data were changed from one form to another, and new levels on the blackboard and knowledge sources were added to extend the scope of the task. Each change in these systems required very little work relative to the magnitude of the changes. The robustness of blackboard systems lies primarily in the organization of the systems which tends to localize changes. Appropriate modularization of the solution space and knowledge into modules that require little intermodule communication ensures that changes are locally confined. Changes in the problem-solving behavior (knowledge-application strategy) are confined to the control component. Additions and modifications to the knowledge base involve additions of new knowledge sources or are confined to additions and modifications within the existing knowledge sources.

The blackboard organization is also being used as a development tool for many programs for several closely related reasons:[65] First, the blackboard is used primarily as a means of communication between disparate processing modules. The information being communicated is based on the task semantics. Thus, if a signal processor produces a radar track in terms of a vector and a covariance, the information is placed on the blackboard in terms of coordinates, a heading, a speed, and an error estimate. This form of information is process and data-structure independent and can be used by other knowledge sources.

Second, the blackboard approach allows the postponement of design decisions. For example, in developing a complex, robot navigation system with multiple sensors, a system is needed that allows experimentation with different sensors which interact differently with each other. Until an appropriate combination of sensors is found, the system needs to be open ended. Third, the blackboard approach provides freedom from message-passing constraints. The message-passing paradigm requires a recipient of a message as well as a sender. Often, the recipient is not known, or the recipient might have been deleted. In the blackboard approach, the message-sending module places the information on the blackboard, and the developer of the module is

---

[65]This is a partial summary of discussions held at the Blackboard Workshop on 12-13 June, 1986 at Carnegie-Mellon University. The attendees were primarily representatives from vision and robotics projects. It appears that for these projects at this point in time the blackboard approach is being used primarily as a software engineering tool.

freed from worrying about other modules.

Not all of this is good news. If the initial organization of the blackboard data (and indirectly the organization of the knowledge sources) is wrong, then modifications result in a rapid deterioration of the system structure. Where can things go wrong? Given a task domain there are many ways to partition its solution space into a hierarchy. If the task (a goal to be accomplished within the domain) is misunderstood, then one can end up with an inappropriate hierarchy on the blackboard. A knowledge engineer will be building a system with a wrong model of the world. Unfortunately, it often takes a long time to discover that one's perspective on the problem is wrong. To aggravate the situation, a clever programmer can go a long way to make things work, even knowing that things will get progressively worse as more knowledge is added. In such a situation, changes occur too late, making a major overhaul inevitable. This type of headache can sometimes be avoided by employing an incremental development strategy, which is easily accomplished with the blackboard system organization.

Blackboard systems can be built in a top-down fashion similar to a standard software engineering technique. A small amount of knowledge can be encoded for each knowledge source, and a simple control component can be constructed to test the overall behavior of the system. This is particularly easy to do if the knowledge sources contain rules because it is easy to implement a few rules for each knowledge source. This approach is in contrast to an approach in which one knowledge source is completed before the next knowledge source is developed.[66] An advantage of this top-down approach is that gross errors, misunderstandings, or inadequacies in the implementation can be discovered quite early. The approach is analogous to the progression of prototypes leading to a fieldable system that was discussed earlier. The increments of development are much smaller, and it permits better internal control of the development process.

4.3. The Blackboard Model as a Research Tool

The blackboard model can be a useful tool for conducting research in applied artificial intelligence. The solution space of the application problem and the domain knowledge can be partitioned in many ways, and a variety of reasoning strategies can be experimented with. The blackboard model itself is also the focus of some current research projects:

- BB-1: The work on the BB-1 skeletal system addresses the problem of rationalizing

---

[66]I have often seen development of a knowledge source that proved unnecessary or unimportant when it was used in conjunction with other knowledge sources. One might argue that a wrong system decomposition was chosen or that the interactions among the subsystems were not fully explored. More often than not, however, the perspective on the problem and on the problem-solving strategy change once all the knowledge sources are in place.

the control component.- The problem of when and how to apply domain knowledge
is viewed as a separate task. This control task is organized as a separate blackboard
system. To solve a problem, BB-1 alternates between formulating and executing the
problem-solving strategy on one blackboard system and applying the domain
knowledge in another. [18]

- **Distributed Problem Solving:** This research addresses the problems of obtaining a
solution and maintaining coherent problem-solving behavior in distributed systems
that have a common goal. One of the major research issues in distributed
knowledge-based systems is the distribution of control -- how much and what kinds
of global control are required and how much control can remain local. This and
other issues are explored within the context of three application domains:
distributed interpretation, distributed network traffic light control, and distributed
planning. [26]

- **Concurrent Problem Solving:** This research explores the feasibility of using the
blackboard model as a basis for developing frameworks for concurrent problem
solving. Two experimental frameworks are being developed to explore parallel
constructs, one for shared-memory, multiprocessor systems and one for distributed-
memory, multiprocessor systems. Major issues are finding and expressing
parallelism in application problems; determining the optimal grain-size of data and
knowledge for maximal processing speed; and, as in the distributed problem-solving
problem, determining the balance between local and global controls. [37]

In addition to this research, it might be interesting and useful to explore the utility of the
blackboard model in the area of design and machine learning. Designing, whether designing a
piece of jewelry or a very large scae interpgration (VLSI) circuit, is an opportunistic process
and is accomplished at various levels of abstraction. Design problems also have large solution
spaces and require many, diverse cooperating sources of knowledge. Design seems a good
candidate for a blackboard application. The major difficulty with design problems might lie
in our limited understanding of how to represent and reason about spatial relationships.

In the case of machine learning, the blackboard model might serve as a model of learning.
The learning process can be viewed as incremental (acquiring a piece of knowledge at a time),
opportunistic (recognizing when something is worth remembering), and hierarchical (knowledge
ranging from detailed empirical association rules to general rules). A learning system might
consist of two blackboard systems. One system, which solves a problem, would consist of the
standard blackboard configuration. The second system, which learns from the behavior of the
first system, would consist of a blackboard containing a hypothesis about how the problem is
being solved, and its knowledge sources would consist of diverse methods of learning. Whether
the learning subsystem performed well or not can be tested by executing the acquired
knowledge on the problem-solving blackboard. This particular approach would require

extensions to the current blackboard organization. These extentions pose little difficulty, however, because the power of the blackboard system organization lies in its modularity, flexibility, and robustness.

Acknowledgements

References

[1]   Nelleke Aiello.
        A Comparative Study of Control Strategies for Expert System: AGE Implementation of
            Three Variations of PUFF.
        Proceedings of the National Conference on Artificial Intelligence :1 - 4, 1983.

[2]   Nelleke Aiello.
        User-Directed Control of Parallelism: The CAGE System.
        Technical Report KSL Report 86-31, Knowlege Systems Laboratory, Computer Science
            Department, Stanford University, April, 1986.

[3]   Harold Brown, Jack Buckman, et. al.
        Final Report on HANNIBAL.
        Technical Report, ESL, Inc., 1982.
        Internal document.

[4]   William J. Clancey.
        Heuristic Classification.
        Technical Report KSL-85-5, Knowledge Systems Laboratory, Computer Science
            Department, Stanford University, 1985.

[5]   Harold Cohen.
        What Is an Image.
        Proceedings of the 6th International Joint Conference on Artificial Intelligence :1028
            - 1051, 1977.

[6]   Randall Davis and Jonathan King.
        An Overview of Production Systems.
        In E.W. Elcock and D. Michie (editor), *Machine Intelligence 8: Machine
            Representation of Knowledge,* . John Wiley, New York, 1977.

[7]   J. Robert Ensor and John D. Gabbe.
        Transactional Blackboards.
        Proceedings of the 9th International Joint Conference on Artificial Intelligence :340
            - 344, 1985.

[8]   Lee D. Erman, Frederick Hayes-Roth, Victor R. Lesser, D. Raj Reddy.
        The Hearsay-II Speech Understanding System: Integrating Knowledge to Resolve
            Uncertainty.
        ACM Computing Survey 12:213 - 253, June, 1980.

[9]   Lee D. Erman, Phillip E. London, Stephen F. Fickas.
        The Design and an Example Use of HEARSAY-III.
        Proceedings of the 7th International Joint Conference on Artificail Intelligence :409
            - 415, 1981.

[10]   Edward A. Feigenbaum and Julian Feldman. (eds).
        Computers and Thought.
        McGraw-Hill Book Company, New York, New York, 1963".

[11] Edward A. Feigenbaum.
The Art of Artificial Intelligence: I. Themes and Case Studies of Knowledge
    Engineering.
Proceedings of the 5th International Joint Conference on Artificial Intelligence :1014
    - 1029, 1977.

[12] Richard D. Fennell and Victor R. Lesser.
Parallelism in AI Problem Solving: A Case Study of HEARSAY-II.
IEEE Transactions on Computers :98 - 111, February, 1977.

[13] Charles Forgy and John McDermott.
OPS, A Domain-Independent Production System Language.
Proceedings of the 5th International Joint Conference on Artificial Intelligence
    1:933-939, 1977.

[14] Charles Forgy, Anoop Gupta, Allen Newll and Robert Wodig.
Initial Assessment of Architectures for Production Systems.
Proceedings of the National Conference on Artificial Intelligence :116 - 119, 1984.

[15] Frederick Hayes-Roth and Victor R. Lesser.
Focus of Attention in a Distributed Logic Speech Understanding System.
Technical Report, Computer Science Department, Carnegie-Mellon University, January,
    1976.

[16] Barbara Hayes-Roth.
The Blackboard Architecture: A General Framework for Problem Solving?.
Technical Report HPP-83-30, Knowledge Systems Laboratory, Computer Science
    Department, Stanford University, May, 1983.

[17] Frederick Hayes-Roth, Donald Waterman, and Douglas Lenat (eds).
Building Expert Systems.
Addison-Wesley, 1983.

[18] Barbara Hayes-Roth.
Blackboard Architecture for Control.
Journal of Artificial Intelligence 26:251 - 321, 1985.

[19] Barbara Hayes-Roth, et al.
Elucidating Protein Structure from Caonstraints in PROTEAN.
In Robert Engelmore and Anthony Morgan (editor), *Blackboard Systems: Applications
    and Frameworks,* . Addison-Wesley, forthcoming.

[20] Barbara Hayes-Roth, Frederick Hayes-Roth, Stan Rosenschein, and Stephanie Cammarata.
Modelling Planning as an Incremental, Opportunistic Process.
Proceedings of the 6th International Joint Conference on Artificial Intelligence
    :375-383, 1979.

[21] Paul J. Kline and Steven B. Dolins.
Choosing Architectures for Expert Systems.
Technical Report RADC-TR-85-192, Texas Instruments Inc., October, 1985.

[22] John Kunz, R. Fallat, D. McClung, J. Osborn, B. Votteri, H.P. Nii, J.S. Aikins, L. Fagen, and E.A. Feigenbaum.
A Physiological Rules Based System for Interpreting Pulmonary Function Test Results.
Technical Report HPP-78-19, Knowledge Systems Laboratory, Computer Science Department, Stanford University, December, 1978.

[23] W. L. Lakin and J. A. H. Miles.
A Blackboard System for Multi-Sensor Fusion.
Technical Report, ASWE, Portsdown, Portsmouth PO6 4AA, England, February, 1984.

[24] Victor R. Lesser, Richard D. Fennell, Lee D. Erman, and D. Raj Reddy.
Organization of the HEARSAY-II Speech Understanding System.
In *IEEE Symposium on Speech Recognition, Contributed Papers*, pages 11-M2 - 21-M2. IEEE Group on Acoustics, Speech and Signal Processing, Computer Science Department, Carnegie-Mellon University, April, 1974.

[25] Victor R. Lesser and Lee D. Erman.
The Retrospective View of the HEARSAY-II Architecture.
Proceedings of the 5th International Joint Conference on Artificial Intelligence :790 - 800, 1977.

[26] Victor R. Lesser and Daniel D. Corkill.
The Distributed Vehicle Monitoring Testbed: A Tool for Investigation Distributed Problem Solving Networks.
The AI Magazine Fall:15 - 33, 1983.

[27] Robert Lindsay, Bruce G. Buchanan, Eeward A. Feigenbaum, and Joshua Lederberg.
Applications of Artificial Intelligence for Organic Chemistry: The Dendral Project.
McGraw-Hill, New York, 1980.

[28] J.L. McClelland and D.E. Rumelhart.
An interactive Activation Model of Context effects in letter Perception: Part 1, An Account of Basic Findings.
Psychological Review 88:375 - 407, 1981.

[29] Brian P. McCune and Robert J. Drazovich.
Radar with Sight and Knowledge.
Defense Electronics , August, 1983.

[30] Makoto Nagao, and Takashi Matsuyama.
A Structural Analysis of Complex Aerial Photographs.
Plenum Press, New York, 1980.

[31] Allan Newell.
Heuristic Programming: Ill-Structured Problems.
In J. Aronofsky (editor), *Progress in Operations Research*, pages 360 -414. Wiley, New York, 1969.

[32] Allen Newell and Herbert A. Simon.
Human Problem Solving.
Prentice Hall, Englewood Cliffs, N.J., 1972.

[33] A. Newell, J. Barnett, C. Green, D. Klatt, J. C. R. Licklider, J. Munson, R. Reddy, and
W. Woods.
Speech Understanding System: A Final Report of a Study Group.
North-Holland, 1973.

[34] H. P. Nii, and E. A. Feigenbaum.
Rule-Based Understanding of Signals.
In D.A. Waterman and R. Hayes-Roth (editors), *Pattern-Directed Inference Systems*,
pages 483 - 501. Acedemic Press, 1978.

[35] H. Penny Nii and Nelleke Aiello.
AGE: A Knowledge-based Program for Building Knowledge-based Programs.
Proc. of IJCAI 6 :645 - 655, 1979.

[36] H. Penny Nii, Edward A. Feignebaum, John J. Anton, and A. Joseph Rockmore.
Signal-to-Symbol Transformation: HASP/SIAP Case Study.
AI Magazine 3:2:23 - 35, 1982.

[37] H. Penny Nii.
*CAGE and POLIGON: Two Frameworks for Blackboard-based Concurrent Problem
Solving.*
Technical Report KSL-86-41, Knowledge Systems Laboratory, Computer Science
Department, Stanform University, 1986.
also in the Proceedings of Expert Systems Workshop.

[38] D. Raj Reddy, Lee D. Erman, and Richard B. Neely.
A Model and a System for Machine Recognition of Speech.
IEEE Transactions on Audio and Electroacoustics AU-21:229 - 238, June, 1973.

[39] D. Raj Reddy, Lee D. Erman, and Richard B. Neely.
The HEARSAY Speech Understanding System: An Example of the Recognition Process.
Proceedings of the 3rd International Joint Conference on Artificial Intelligence :185
- 193, 1973.

[40] James Rice.
Poligon: A System for Parallel Problem Solving.
Technical Report KSL-86-19, Knowledge Systems Laboratory, Computer Science
Department, Stanford University, April, 1986.

[41] D.E. Rumelhart and J. L. McClelland.
An Interactive Model of Context Effects in Letter Perception: Part 2, The Enhancement
Effect and Some Tests and Extensions to the Model.
Psychological Review 89:60 - 94, 1982.

[42] Oliver G. Selfridge.
Pandamonium: A Paradigm for Learning.
Proceedings of the Symposium on the Mechanization of Thought Processes :511 - 529,
1959.

[43] Steven A. Shafer, Anthony Stentz, and Charles Thorpe.
An Architecture for Sensor Fusion in a Mobile Robot.
Proceedings of the 1986 IEEE International Conference on Robotics and Automation ,
1986.
To appear in the proceedings.

[44] Beau Sheil.
      Power Tools for Programmers.
      Datamation :131 - 144, February, 19~ .

[45] Edward H. Shortliffe.
      Computer-Based Medical Consultation: MYCIN.
      American Elsevier, New York, 1976.

[46] Herbert A. Simon.
      The Sciences of the Artificial.
      Massachusetts Institute Technology Press, Cambridge, Mass, 1969.

[47] Herbert A. Simon.
      Scientific Discovery and the Psychology of Problem Solving.
      In *Models of Discovery, .* D. Reidel Publishing Company, Boston, Mass, 1977.

[48] David S. Spain.
      Application of Artificial Intelligence to Tactical Situation Assessment.
      Proceedings of the 16th EASCON83 :457 - 464, September, 1983.

[49] Anthony Stentz and Steve Shafer.
      Module Programmer's Guide to Local Map Builder for ALVan.
      Technical Report, Carnegie-Mellon University Computer Science Departement, August,
          1985.

[50] Allan Terry.
      The CRYSALIS Project: Hierachical Control of Production Systems.
      Technical Report HPP-83-19, Stanford University, Heuristic Programming Project, May,
          1983.

[51] Donald A. Waterman.
      A Guide to Expert Systems.
      Addison-Wesley, 1985.

[52] Sholom Weiss and Casimir Kulikowski.
      A Practical Guide to Building Expert Systems.
      Rowman & Allanheld, New Jersey, 1984.

[53] Mark Williams, Harold Brown, and Terry Barnes.
      TRICERO Design Description.
      Technical Report ESL-NS539, ESL, Inc., May, 1984.

[54] Mark A. Williams.
      Distributed, Cooperating Expert Systems for Signal Understanding.
      Proceedings of Seminar on AI Applications to Battlefield :3.4-1 - 3.4-6, 1985.

END

12-86

DTIC